arXiv:1401.6196v1 [cs.CV] 23 Jan 2014

SPATIALLY REGULARIZED RECONSTRUCTION OF FIBRE ORIENTATION DISTRIBUTIONS IN THE PRESENCE OF ISOTROPIC DIFFUSION

Q. ZHOU, O. MICHAILOVICH, AND Y. RATHI

Abstract. The connectivity and structural integrity of the white matter of the brain is nowadays known to be implicated into a wide range of brain-related disorders. However, it was not before the advent of diffusion Magnetic Resonance Imaging (dMRI) that researches have been able to examine the properties of white matter *in vivo*. Presently, among a range of various methods of dMRI, high angular resolution diffusion imaging (HARDI) is known to excel in its ability to provide reliable information about the local orientations of neural fasciculi (aka fibre tracts). Moreover, as opposed to the more traditional diffusion tensor imaging (DTI), HARDI is capable of distinguishing the orientations of multiple fibres passing through a given spatial voxel. Unfortunately, the ability of HARDI to discriminate between neural fibres that cross each other at acute angles is always limited, which is the main reason behind the development of numerous post-processing tools, aiming at the improvement of the directional resolution of HARDI. Among such tools is spherical deconvolution (SD). Due to its ill-posed nature, however, SD standardly relies on a number of *a priori* assumptions which are to render its results unique and stable. In this paper, we propose a different approach to the problem of SD in HARDI, which accounts for the spatial continuity of neural fibres as well as the presence of isotropic diffusion. Subsequently, we demonstrate how the proposed solution can be used to successfully overcome the effect of partial voluming, while preserving the spatial coherency of cerebral diffusion at moderate-to-severe noise levels. In a series of both *in silico* and *in vivo* experiments, the performance of the proposed method is compared with that of several available alternatives, with the comparative results clearly supporting the viability and usefulness of our approach.

Key words: diffusion imaging, MRI, spherical deconvolution, HARDI, sparse analysis, fibre continuity, and total variation

1. Introduction

The connectivity and structural integrity of white matter is nowadays known to be indicative of a wide range of brain-related pathologies. While "invisible" to alternative means of imaging-based diagnosis, the above information can be elicited from the measurements acquired by means of *diffusion Magnetic Resonance Imaging* (dMRI). This fact has triggered an active development of various dMRI methodologies, which has made dMRI into a well-established technique of modern medical imaging [1].

At the present time, dMRI encompasses a number of various methodologies and protocols, the most widely acknowledged of which is Diffusion Tensor Imaging (DTI) [2, 3, 4, 5, 6].

It is known, however, that the modelling capacity of DTI is limited due to its reliance on assuming the ensemble averaged diffusion propagator (EAP) to be a unimodal Gaussian. In fact, the above assumption undermines the ability of DTI to provide accurate estimation of the apparent diffusivity of white matter at the locations of crossing, diverging, and kissing neural fibre tracts [7, 8, 9]. Alternatively, a parameter-free approach to estimate EAP is provided by Diffusion Spectral Imaging (DSI) [10], which allows reconstruction of complex diffusivity profiles under rather general conditions. Unfortunately, practical implementation of DSI entails acquisition of diffusion measurements over a dense Cartesian grid in the q-space, which renders the acquisition requirements of DSI beyond the limits of practically admissible. This problem, however, can be alleviated by restricting the diffusion measurements to a relatively small number of concentric shells in the q-space. This sampling strategy – known as Multi-Shell Diffusion Imaging (MSDI) [11, 12, 13] – has successfully served as a basis for many advanced dMRI methodologies [11, 14, 15, 16, 17, 18, 19, 20].

Although the availability of the EAP is generally preferred, in some cases it is sufficient to know the result of its marginalization over the range variable. The resulting probability density is known as the orientation distribution function (ODF), and it quantifies the probability with which water molecules undergo displacement along various spatial directions [21, 22]. On the practical side, a useful approximation of the ODF can be obtained by means of q-ball imaging (QBI) [22], which can in turn be based on a single-shell data acquisition scheme, known as High Angular Resolution Diffusion Imaging (HARDI) [23, 24, 25, 26, 8, 27, 28].

The nature of single-shell acquisition inherent in HARDI imposes constraints on the directional resolution of estimated ODFs, with higher values of the diffusion scintillation parameter b (the so-called b-value) leading to better resolvability between various diffusion modes within a given voxel. This fact is exemplified in Fig. 1, which shows a simulated ODF[1] (top subplot) corresponding to two fibre tracts crossing each other at an angle of 60^{o}. At the same time, the bottom row of subplots depict the ODFs which have been recovered from the associated HARDI data generated with $b \in \{1000, 3000, 5000\}$ s/mm^2. One can see that the best directional resolution is attained at the maximum value of b= 5000 s/mm^2, as expected. It is also worthwhile noting that the above effect is intrinsic in both the Funk-Radon transform (FRT) [22] and solid angle [27, 29] formulations of QBI.

As tempting as it might seem at the first glance, working with relatively high b-values is usually avoided in practice for two main reasons. First, an increase in the b-value is typically achieved through using longer mixing times, which unavoidably leads to substantially low values of signal-to-noise ratio (SNR) [30]. Second, using higher b-values makes the diffusion data less sensitive to the effects of fast diffusion, which is often associated with the concept of "free water" [31], and hence represents a diagnostically important diffusion regime. As a result, diffusion data are commonly acquired with relatively low values of b, typically around 1000 s/mm^2. Needless to add, in such cases, poor directional resolution of ODFs may become an issue, especially for such applications as fibre tractography [32].

[1]The ODF was generated using a standard Gaussian mixture model [23] with equal volume fractions, FA=0.8 and MD = $7 \cdot 10^{-4}$ mm^2/s.

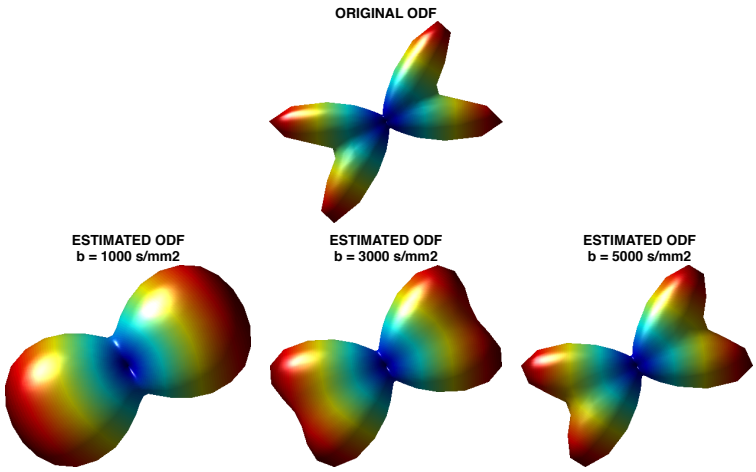

FIGURE 1. (Top) Original ODF; (Bottom) Estimated ODFs obtained from HARDI data generated with (left to right) b=1000, 3000, and 5000 s/mm^2.

The problem of limited directional resolution of QBI can be addressed using the framework of spherical deconvolution (SD) [33]. In this formulation, SD is used to recover a *fibre ODF* (fODF) which, as opposed to ODF, quantifies the likelihood of a neural fibre to have a (local) tangent vector of a certain orientation. The idea of using SD as a method for improving the directional resolution of QBI was first introduced in [34]. In this initial approach, fODFs were recovered through a direct matrix inversion (aka naive deconvolution) preceded by low-pass filtering. A more advanced SD technique based on a least-square (LS) formulation with positivity constraints was proposed by the same authors in [35, 36]. An interesting extension to the later method was also discussed in [37], where the uncertainty in fODF estimation was reduced through a bootstrap procedure. In [38], the fOFD is parameterized based on a two-compartment model, followed by its reconstruction using non-regularized (naive) inverse filtering. A similar line of arguments is used in [39], albeit this time with Tikhonov regularization implemented implicitly through damped SVD-based inversion. Moreover, as opposed to many other works in the field, this study addressed the important problem of automatic determination of an optimal value of the regularization parameter based on generalized cross validation. Advanced statistical considerations were also employed in [40] to derive a maximum entropy deconvolution algorithm. The same idea of entropy maximization was later adopted by [41]. In the latter work, however, the reconstruction was performed within the framework of blind source separation, which offers the substantial advantage of independence on the knowledge of single fibre response. Unfortunately, the Infomax approach used by the authors in [41] does not explicitly take into account the distribution and level of measurement noises, which makes it overly dependent on measurement conditions. Finally, we would also like to mention the SD approach of [42], which is close in philosophy to the methods discussed in the present paper. In particular,

[42] advocates the idea of *sparse* SD, congruent with the assumption on an fODF to have a relatively small number of "sharp" maxima in the directions of associated neural bundles.

While different in their assumptions, computational requirements, and performances, virtually all the above-mentioned SD methods share a common drawback. Particularly, all these methods strongly adhere to the assumption on the data signal to be formed as a convolution of a *single-fibre response* (SRF) and an fODF, thereby explicitly requiring the presence of a fibre tract at the corresponding voxel. However, situations are frequent in which the data may alternatively be associated with either grey matter or cerebrospinal fluid (CSF), as well as with partial volume fraction voxels containing a mixture of different types of cerebral tissue. In such cases, the diffusion signals should be expected to have a non-negligible *isotropic* component, which comes at odds with the assumption on fODFs to be, e.g., sparse [42] or of maximal entropy [40].

The above-mentioned deficiency of earlier application of SD to dMRI has been alleviated in a series of more recent works on the subject. Thus, in [43, 44], fODF has been supplanted by a different quantity, called a "fibre orientation function" (FOF), which explains a combined effect of both (multimodal) anisotropic and isotropic diffusions. Subsequently, a variation of the Lucy-Richardson deconvolution algorithm was employed to recover the fODF, while gradually suppressing the isotropic component of the signal. The deconvolution algorithm in [45] iteratively estimates and subtracts the isotropic component from HARDI data before fODFs are estimated. A data formation model similar to that of [43, 44] was also exployed by the authors of [46], in combination with the sparse deconvolution ideas of [42]. Quite a different type of deconvolution methods were explored in [47, 48] based on Bayesian inference. Specifically, [48] takes advantage of a measure-theoretic framework to represent the fODF as a probability measure decomposable into three unique components, which allow adequate description of a wide spectrum of possible fODF shapes and patterns. Similarly to the previously mentioned results, the models used in [47, 48] explicitly account for the presence of isotropic diffusion.

Since neutral bundles extend continuously in space, it is reasonable to assume spatially adjacent fODFs to exhibit a fair amount of correlation, which can be exploited to improve the results of SD. This idea has been exploited in several studies to improve the accuracy of QBI [49, 50]. In application to SD, spatial regularization was used in [51], in which case a spatial continuity of fODFs was enforced via minimizing a weighted quadratic penalty. A more recent work in [52] has introduced an anisotropic regularization scheme which, for any spatial direction $\mathbf{u}$, minimizes the $\mathcal{L}_2$-norm of the projection of the spatial gradient of fODFs onto $\mathbf{u}$. It has been shown that this minimization favours the spatial continuity of neural fibres, thereby producing anatomically plausible reconstructions. Unfortunately, neither of the above-mentioned approaches was designed to deconvolve HARDI signals in the presence of isotropic diffusion.

Despite the apparent success of SD in application to diffusion imaging, there is still much space for further improvements. In particular, it appears that little has been done on the development of SD algorithms which can perform reliably in the presence of isotropic diffusion, while imposing effective regularization constraints on *both* the isotropic and anisotropic components of diffusion signals. Needless to add, contriving such a reconstruction problem

is likely to result in composite optimization, solving which could be a non-trivial problem by itself. Accordingly, the present paper aims to contribute to the existing body of works on SD in the following directions:

(1) We formulate a new method for SD of HARDI data, subject to spatial regularization of *both* the isotropic and anisotropic components of HARDI signals as well as their related ODFs/fODFs.
(2) We describe a computationally efficient implementation of the proposed algorithm based on the idea of variable splitting [53]. The proposed computational solution has a particularly simple modular structure, which is straightforward to reproduce using standard computational means.

The remainder of the paper is organized as follows. Section 2 details the data formation model used in the subsequent derivations of the proposed deconvolution approach, which is introduced in Section 3. Section 4 provides details on our numerical implementation of the proposed method, while Section 5 sets out the experimental setup used for its numerical validation. The results of application of the proposed method to both computer simulated and *in vivo* data are reported in Section 6. The section also contains a comparative analysis of the performance of the proposed algorithm against that of several alternative solutions. Finally, Section 7 finalizes the paper with recapitulation of its results and conclusions.

2. Data formation model

To fix the ideas, we start with a formal setting, in which HARDI data are assumed to be collected over a bounded (open) subset of $\Omega \in \mathbb{R}^3$. In this case, for each $\mathbf{r} \in \Omega$, a HARDI signal $s(\mathbf{u}|\mathbf{r})$ can be viewed as a positive-valued, spherical function $s(\cdot|\mathbf{r}) : \mathbb{S}^2 \to [0, \infty)$, with $\mathbb{S}^2$ denoting the unit sphere in $\mathbb{R}^3$, in which case $\mathbf{u} \in \mathbb{S}^2$ is interpreted as the direction of diffusion encoding. When a given voxel of interest supports a single neural fibre, the corresponding EAP can be closely approximated by a unimodal Gaussian density, in which case the HARDI signal $s(\mathbf{u}|\mathbf{r})$ can be described as [2, 4]

$$s(\mathbf{u}|\mathbf{r}) = s_0(\mathbf{r}) \exp\{-b\ \mathbf{u}^T D(\mathbf{r})\mathbf{u}\}, \quad \forall \mathbf{r} \in \Omega, \tag{1}$$

where the diffusion tensor $D(\mathbf{r})$ encodes the directivity and ellipticity of the pattern of local diffusion, while the b-value is typically set in the range between 1000 and 3000 s/mm^2. Note that the b_0-image $s_0(\mathbf{r})$ is usually acquired through additional measurements and used to normalize the HARDI signals. For the sake of notational convenience, in what follows, the signal $s(\mathbf{u}|\mathbf{r})$ will be assumed to be normalized, implying $s_0(\mathbf{r}) = 1$, for all $\mathbf{r} \in \Omega$.

Although standard in DTI, the model of (1) is not applicable in situations when a voxel of interest supports multiple neural bundles [9]. If this is the case, then under some fairly general conditions [23], the HARDI signal $s(\mathbf{u}|\mathbf{r})$ can be assumed to be formed as a (linear) superposition of several "DTI signals" weighted by their respective partial volume fraction coefficients. To formalize such a signal formation model, it is common to use the notion of an SFR [34], which is, in fact, equal to the elementary "DTI signal" given by (1). Moreover, in the SD literature, for the definition of a SFR $h(\mathbf{u})$ it is standard to use a cylindrically symmetric diffusion tensor $D_0 = \mathrm{diag}\{\lambda_\perp,\ \lambda_\perp,\ \lambda_\parallel\}$ (with $\lambda_\parallel > \lambda_\perp$), which

seems to naturally comply with an expected (local) geometry of neural fasciculi. In such a case, it is straightforward to show that the SFR becomes a *zonal* (spherical) function [54] that is formally given by

$$h_{\mathbf{v}_0}(\mathbf{u}) = h(\mathbf{u}\cdot\mathbf{v}_0) = \alpha\exp\{-\beta(\mathbf{u}\cdot\mathbf{v}_0)\}, \tag{2}$$

with the dot standing for the standard (Euclidean) dot product in $\mathbb{R}^3$, $\alpha \triangleq \exp\{-b\lambda_\perp\}$, $\beta \triangleq b(\lambda_\parallel - \lambda_\perp)$, and $\mathbf{v}_0 = [0,\ 0,\ 1]^T$ denoting the north pole of $\mathbb{S}^2$.

It is worthwhile noting that, due to the property of the SFR in (2) to be zonal, its value at a given $\mathbf{u}$ depends only on the angle between $\mathbf{u}$ and a fixed direction $\mathbf{v}$ (e.g., $\mathbf{v} = \mathbf{v}_0$, as in (2)), which makes $h_{\mathbf{v}}(\mathbf{u})$ invariant under rotations around $\mathbf{v}$. This invariance allows the result of convolution with $h_{\mathbf{v}}(\mathbf{u})$ to be expressed as a function of $\mathbb{S}^2$ (rather than of the orthogonal group $SO(3)$ [33]), in which case our signal formation model becomes

$$s(\mathbf{u}|\mathbf{r}) = \int_{\mathbb{S}^2} h_{\mathbf{u}}(\mathbf{v})d\mu(\mathbf{v}|\mathbf{r}) = \int_{\mathbb{S}^2} h(\mathbf{u}\cdot\mathbf{v})d\mu(\mathbf{v}|\mathbf{r}), \quad \forall\mathbf{r}\in\Omega. \tag{3}$$

Here $\mu(\mathbf{u}|\mathbf{r})$ is a probability measure that is used to model the fibre probability distribution over $\mathbb{S}^2$ [48]. In particular, at any $\mathbf{r}\in\Omega$, $\mu(\cdot|\mathbf{r}) : \mathcal{B} \to [0,\infty]$ quantifies the relative frequency of specific fibre orientations over a given element of the Borel sigma algebra $\mathcal{B}$ of $\mathbb{S}^2$.

A mathematically elegant and physiologically meaningful way to interpret the structure of μ was recently described in [48], where the authors took advantage of the Lebesgue's decomposition theorem to represent μ as a sum of three components (*viz.*, discrete, absolutely continuous, and singular continuous), each of which is able to model a distinct characteristic of the fibre orientation distribution. In the current paper, however, we proceed under a simplified assumption on μ to be absolutely continuous, in which case it can be described in terms of a non-negative, Borel measurable function $f(\mathbf{u}|\mathbf{r})$ as $d\mu(\mathbf{v}|\mathbf{r}) = f(\mathbf{v}|\mathbf{r})d\eta(\mathbf{v})$, with η being the Haar measure of $\mathbb{S}^2$. We note that this simplification seems to be reasonable, considering the fact that, in practical computations, both measurements and resulting estimates are always bounded in value and discrete.

The above simplifying assumption leads to the standard (forward) model for $s(\mathbf{u}|\mathbf{r})$ which reads [36, 42, 44, 48, 46]

$$s(\mathbf{u}|\mathbf{r}) = \int_{\mathbb{S}^2} h(\mathbf{u}\cdot\mathbf{v})f(\mathbf{v}|\mathbf{r})d\eta(\mathbf{v}), \quad \forall\mathbf{r}\in\Omega, \tag{4}$$

in which case the density $f(\mathbf{u}|\mathbf{r})$ is conventionally referred to as a *fibre orientation distribution function* (fODF). It should also be noted that the models in (3) and (4) are *stationary*, since the SFR is assumed to be fixed within a given voxel as well as across the whole image domain Ω. While only approximative [55], the stationary SD model has nevertheless shown to yield useful reconstructions, while offering the important advantage of tractability and amenability to numerical computations.

Additionally, the structure of fODF $f(\mathbf{u}|\mathbf{r})$ in (4) deserves a special consideration. Since the measured diffusion signal receives contributions from both coherently ordered axonal fascicles as well as from their complex and more heterogeneous extra-axonal surroundings (containing astrocytes, glia, and randomly oriented extracellular matrix molecules) [16], it

seems reasonable to consider $f(\mathbf{u}|\mathbf{r})$ to be composed of two main terms, *viz.* anisotropic and isotropic. Specifically, following the line of ideas advocated in [44, 48, 46], we model $f(\mathbf{u}|\mathbf{r})$ according to

$$f(\mathbf{u}|\mathbf{r}) = (1 - p_{\text{iso}}(\mathbf{r}))f_a(\mathbf{u}|\mathbf{r}) + p_{\text{iso}}(\mathbf{r})f_{\text{iso}}(\mathbf{r}), \tag{5}$$

with $f_a(\mathbf{u}|\mathbf{r})$ and $f_{\text{iso}}(\mathbf{r})$ representing the anisotropic and isotropic components of the fibre probability distribution, respectively, and $0 \leq p_{\text{iso}}(\mathbf{r}) \leq 1$ controlling their partial volume fractions at $\mathbf{r} \in \Omega$. Moreover, $f_{\text{iso}}(\mathbf{r})$ can be further represented as $f_{\text{iso}}(\mathbf{r}) = \exp(-b\lambda_{\text{iso}}(\mathbf{r}))$, where $\lambda_{\text{iso}}(\mathbf{r}) > 0$ is the apparent diffusivity of the isotropic component [48, 46].

Since the isotropic components $f_{\text{iso}}(\mathbf{r})$ in (5) is dissociated from the anisotropic component $f_a(\cdot|\mathbf{r})$, it would no longer be correct to regard $f(\mathbf{u}|\mathbf{r})$ as an fODF. To overcome this notational inconsistence, from now on, we will use this term when referring to $f_a(\mathbf{u}|\mathbf{r})$ instead. At the same time, since $f_{\text{iso}}(\mathbf{r})$ can be viewed as a scalar-valued function of $\mathbf{r} \in \Omega$, it will be referred hereinafter to as an *isotropic diffusion map* (IDM).

Finally, in this paper, instead of trying to recover $f_a(\mathbf{u}|\mathbf{r})$ and $f_{\text{iso}}(\mathbf{r})$ along with $p_{\text{iso}}(\mathbf{r})$, we estimate the weighted quantities $f'_a(\mathbf{u}|\mathbf{r}) \triangleq (1-p_{\text{iso}}(\mathbf{r}))f_a(\mathbf{u}|\mathbf{r})$ and $f'_{\text{iso}}(\mathbf{r}) \triangleq p_{\text{iso}}(\mathbf{r})f_{\text{iso}}(\mathbf{r})$, such that $f(\mathbf{u}|\mathbf{r}) = f'_a(\mathbf{u}|\mathbf{r}) + f'_{\text{iso}}(\mathbf{r})$, $\forall \mathbf{r}$. We note that, when normalized by $\int_{\mathbb{S}^2} f(\mathbf{u}|\mathbf{r})d\eta(\mathbf{u})$, $f'_a(\mathbf{u}|\mathbf{r})$ and $f'_{\text{iso}}(\mathbf{r})$ acquire the "flavour" of posterior probabilities. Thus, for example, the values of $f'_a(\mathbf{u}|\mathbf{r})$ are not only indicative of the orientations of neural fibres, but also reflect one's *level of confidence* that the fibres are *actually* present at a given location $\mathbf{r}$ at the first place. We support this concept through an experimental study and argue that using the weighted densities $f'_a(\mathbf{u}|\mathbf{r})$ and $f'_{\text{iso}}(\mathbf{r})$ can benefit a number of related applications, such as probabilistic fibre tractography [32].

3. Proposed approach

3.1. **Model discretization.** As usual, the formalism of matrix-vector multiplications turns out to be the most convenient for formulation of practical solutions. To this end, we first note that (normalized) HARDI data typically consists of a set of diffusion-weighted scans acquired for K directions $\{\mathbf{u}_k\}_{k=1}^K$ of diffusion-encoding gradients (with $\mathbf{u}_k \in \mathbb{S}^2, \forall k$). For the convenience of exposition, we concatenate these scans into K *row* vectors of length I, with I being equal to the number of spatial samples (voxels) within Ω. Further, the vectors thus obtained can be organized as the rows of a $K \times I$ matrix s, in which case the columns of s correspond to diffusion measurements observed at different spatial locations. In what follows, we use both super- and subscripts to distinguish between the columns $\{s^i\}_{i=1}^I$ and rows $\{s_k\}_{k=1}^K$ of the data matrix s.

Discretizing the SD model of (4) is the next step. To this end, let $\{\mathbf{v}_j\}_{j=1}^J$ (with $J > K$) be a set of spherical points over which the values of the fODFs are to be recovered. Then, given estimates[2] $\tilde{\alpha}$ and $\tilde{\beta}$ of the SFR parameters in (2), we define the (k, j)-th element of

[2]To estimate these parameters, it is standard to fit the DTI model of (1) to HARDI signals corresponding to, e.g., corpus callosum. As the latter is predominantly composed of single (commisural) fibre bundles, its related signals can provide a reliable estimate of the SFR, subject to appropriate averaging [34].

a $K \times J$ matrix $H = \{h_{k,j}\}$ according to

$$h_{k,j} = \tilde{\alpha} \exp\{-\tilde{\beta}(\mathbf{u}_k \cdot \mathbf{v}_j)\}. \tag{6}$$

Note that the J columns of resulting H correspond to the SFR rotated in directions $\{\mathbf{v}_j\}_{j=1}^{J}$ and discretized at points $\{\mathbf{u}_k\}_{k=1}^{K}$. Finally, we define Φ to be a $K \times (J+1)$ matrix obtained from H through addition of an extra column of ones, *viz.*

$$\Phi = [H \ \mathbf{1}], \tag{7}$$

where $\mathbf{1} = [1, 1, ..., 1]^T \in \mathbb{R}^K$.

Now, let $\mathbf{r}_i$ be the coordinate of an arbitrary voxel within Ω, with $i = 1, 2, \ldots, I$, and let $f^i \in \mathbb{R}^{J+1}$ be a column vector defined as $f^i = [f'_a(\mathbf{v}_1|\mathbf{r}_i), f'_a(\mathbf{v}_2|\mathbf{r}_i), \ldots, f'_a(\mathbf{v}_J|\mathbf{r}_i), f'_{\text{iso}}(\mathbf{r}_i)]^T$. Then, in the absence of measurement noises and disregarding the effect of discretization, the SD model (4) along with (5) suggest that $s^i = \Phi f^i$, for all $i = 1, 2, \ldots, I$. It is, therefore, convenient to agglomerate all the above model equations into a single one that reads

$$s = \Phi f, \tag{8}$$

where f is a $(J+1) \times I$ matrix, with its columns defined by f^i, with $i = 1, 2, \ldots, I$. Note that the last row f_{J+1} of f is equal to a row-stacked version of the IDM, whilst the first J rows $f_1, \ldots, f_J$ of f can be viewed as row-stacked versions of the images obtained by restricting the fODF $f_a(\mathbf{u}|\mathbf{r})$ to the directions $\mathbf{v}_1, \ldots, \mathbf{v}_J$, respectively.

3.2. **Estimation framework.** For obvious reasons, recovering a useful estimate of f based on model (8) alone is a futile exercise. According to the formalism of Bayesian estimation, to render the reconstruction unique and stable, the model equation needs to be augmented with reasonable *a priori* assumptions on the nature of f in (8). One of such assumptions, which has been proved to be particularly useful for reconstruction of fODFs, is that of *sparsity*. Indeed, the anatomical organization of white matter suggests the number of axonal fascicles running through any given voxel $\mathbf{r}_i$ is likely to be relatively small. This fact, in turn, implies that the vector $[f'_a(\mathbf{v}_1|\mathbf{r}_i), f'_a(\mathbf{v}_2|\mathbf{r}_i), \ldots, f'_a(\mathbf{v}_J|\mathbf{r}_i)]^T$ could be reasonably expected to have a relatively small number of significant components, with the rest of its entries distributed in a close proximity of zero. Such a behaviour of $f_a(\mathbf{u}|\mathbf{r})$ can be modelled in a number of different ways [42, 48, 46]. In the present paper, we take advantage of the standard method of recovering a sparse vector through minimization of its ℓ_1-norm [42].

Before proceeding to the next step, we note that the SD model of (8) admits an alternative interpretation, according to which every s^i is approximated by a linear combination of the columns of Φ. As the last (constant) column of Φ has distinctly different *morphological* properties as compared to the other columns of the matrix, it seems reasonable to minimize the ℓ_1-norm of the *entire* vector f^i (rather than only of its part associated with $f'_a(\mathbf{u}|\mathbf{r}_i)$). In this case, the parsimonious nature of ℓ_1 minimization will force the optimal solution to be dominated by either an IDM or an fODF component, while permitting them both only when there is clear evidence of their concurrent existence. It deserves noting that a similar principle has been used in morphological component analysis to decompose a signal of interest into morphologically distinct components [56, 57].

To avoid unnecessary complications in notations, in the derivations that follow, we use $\|\cdot\|_2$ and $\|\cdot\|_1$ to denote the ℓ_2- and ℓ_1-norms of vectors as well as the analogous "entry-wise" norms of matrices. Thus, for example, the ℓ_2-norm of a HARDI signal s can be expressed in two ways as $\|s\|_2^2 = \sum_{i=1}^{I} \|s^i\|_2^2 = \sum_{k=1}^{K} \|s_k\|_2^2$, while the ℓ_1-norm of f can be expressed as $\|f\|_1 = \sum_{i=1}^{I} \|f^i\|_1 = \sum_{j=1}^{J+1} \|f_j\|_1$. With this notation at hand, we formulate the problem of finding an optimal f^* as

$$f^* = \arg\min_f \|f\|_1 \qquad (9)$$
$$\text{subject to } \|\Phi f - s\|_2^2 \leq \epsilon,\ f \geq 0,$$

where ϵ controls the size of measurement and model errors, whereas the point-wise inequality constraint is added to assure that both the IDM and fODFs are non-negative quantities. Note that the problem in (9) is analogous to the one described in [46], apart from the fact that (9) applies to a whole set of HARDI data, rather than to a single voxel.

To facilitate numerical solution, the problem (9) is usually reformulated in its equivalent, unconstrained (Lagrangian) form as given by

$$f^* = \arg\min_f \left\{ \frac{1}{2}\|\Phi f - s\|_2^2 + \lambda\|f\|_1 + \varphi_{\geq}(f) \right\}, \qquad (10)$$

where $\lambda > 0$ is a user-controlled regularization parameter and $\varphi_{\geq}$ denotes the indicator function of the positive orthant. Specifically, $\varphi_{\geq}(f) = 0$, if *all* entries of f are non-negative, and $\varphi_{\geq}(f) = +\infty$, otherwise. The problem (10) can be solved using a variety of methods of non-smooth optimization (e.g., [58]).

Unfortunately, the solution of (10) could be only suboptimal, as it completely disregards any spatial-domain dependencies between the values of f. To palliate this deficiency, we first note that each *row* of f can be considered to be a discrete image defined over the spatial lattice $\{\mathbf{r}_i\}_{i=1}^{I} \in \Omega$ and stacked into a row vector. Particularly, in this interpretation, the first J rows $\{f_j\}_{j=1}^{J}$ can be viewed as *restrictions* of the fODF to directions $\{\mathbf{v}_j\}_{j=1}^{J}$, while f_{J+1} represents the corresponding IDM. Naturally, the images $\{f_j\}_{j=1}^{J}$ and f_{J+1} have different statistical properties, and therefore they should be regularized in different ways. Thus, to spatially regularize the fODF $f_a(\mathbf{u}|\mathbf{r})$ we adopt the fibre continuity approach of [52], which requires the directional derivative of f_j along $\mathbf{v}_j$ to be relatively small in value. Formally, let $\nabla_d : \mathbb{R}^I \to \mathbb{R}^I$, with $d = 1, 2, 3$, denote the operators of spatial differencing in the direction of x, y, and z coordinates, respectively. Then, for each $j = 1, \ldots, J$, one can then assemble a $3 \times I$ matrix Df_j, with its rows defined by the partial differences of f_j, *viz.* $Df_j = [\nabla_1 f_j^T\ \nabla_2 f_j^T\ \nabla_3 f_j^T]^T \in \mathbb{R}^{3\times I}$. Subsequently, with $\mathbf{v}_j$ being a column vector, the directional derivative of f_j along $\mathbf{v}_j$ is conveniently given by $\mathbf{v}_j^T Df_j$, in which case the approach of [52] calls for minimizing $\sum_{j=1}^{J} \|\mathbf{v}_j^T Df_j\|_2^2$.

It should be noted that the above approach cannot be extended to the IDM f_{J+1}, since the latter is devoid of directional continuity [52]. Yet, for the pure sake of harmonizing the notations, we replace the original minimization of $\sum_{j=1}^{J} \|\mathbf{v}_j^T Df_j\|_2^2$ by minimizing $\sum_{j=1}^{J+1}(1-\delta_{j,J+1})\|\mathbf{v}_j^T Df_j\|_2^2$, where $\delta_{t,r}$ stands for the Kronecker symbol which obeys $\delta_{t,r} = 1$, if $t = r$,

while $\delta_{t,r} = 0$, otherwise. (Note that, since the last component in the above summation is multiplied by a zero weight, the choice of $\mathbf{v}_{J+1}$ is immaterial and it has no effect on the reconstruction procedure as shown later in the paper.) Subsequently, the optimal solution can still be defined as a global minimizer over f, which is now given by

$$f^* = \arg\min_f \left\{ \frac{1}{2}\|\Phi f - s\|_2^2 + \lambda\|f\|_1 + \mu\|f\|_a^2 + \varphi_{\geq}(f) \right\}, \tag{11}$$

where

$$\|f\|_a^2 = \sum_{j=1}^{J+1} (1 - \delta_{j,J+1}) \|\mathbf{v}_j^T D f_j\|_2^2, \tag{12}$$

and $\mu > 0$ is another regularization constant.

Finally, the problem of regularizing the IDM f_{J+1} must not be overseen as well. Since one can reasonably expect f_{J+1} to vary smoothly throughout the brain with the exception of abrupt changes between white matter, grey matter, CSF, as well as the regions of possible brain pathologies, it seems justified to model f_{J+1} as a function of bounded variation (BV) [59]. In particular, the BV model enforces the assumption on f_{J+1} to have a relatively small value of its *total variation* (TV) seminorm $\|f_{J+1}\|_{TV}$ which can be defined as follows. Let $\nabla_1 f_{J+1}[i]$, $\nabla_2 f_{J+1}[i]$, and $\nabla_3 f_{J+1}[i]$ denote the i-th elements of vectors $\nabla_1 f_{J+1}$, $\nabla_2 f_{J+1}$, and $\nabla_3 f_{J+1}$, respectively. Then, the TV seminorm of f_{J+1} can be defined in a standard way as

$$\|f_{J+1}\|_{TV} = \sum_{i=1}^{I} \Big[\sum_{d=1}^{3} |\nabla_d f_{J+1}[i]|^2 \Big]^{1/2}. \tag{13}$$

It is definitely possible to apply the definition of TV to images f_j, with $j = 1, \ldots, J$ as well. However, minimizing these norms would likely mislead the estimation process, since f_j may not be assumed to be piecewise smooth, in general. Still, to balance the notations, we define the TV seminorm of f according to

$$\|f\|_{TV} = \sum_{j=1}^{J+1} \delta_{j,J+1} \|f_j\|_{TV}, \tag{14}$$

which leads to the optimal solution of the form

$$f^* = \arg\min_f \left\{ \frac{1}{2}\|\Phi f - s\|_2^2 + \lambda\|f\|_1 + \mu\|f\|_a^2 + \nu\|f\|_{TV} + \varphi_{\geq}(f) \right\}, \tag{15}$$

with $\nu > 0$ being an additional regularization constant which controls the piecewise smooth behaviour of the IDM. We admit that automatically determining an optimal value of ν, as well as those of λ and μ, is a difficult problem, which extends well beyond the scope of the current paper. It was observed in practice, however, that finding acceptable values of these parameters by wonted trials-and-errors is a much less arduous task, as it might seem at the first glance.

The solution of (15) entails minimizing a non-smooth cost function, which effectively rules out the use of gradient-based methods of numerical optimization. Moreover, the composite nature of the cost makes it difficult to devise an efficient optimization approach which would perform the minimization *directly* with respect to f. To overcome these difficulties, the next section introduces a particularly simple solution using the alternating directions method of multipliers (ADMM) [60]. Apart from breaking down the optimization in (15) into a sequence of simple and closed-form solutions, the method offers a straightforward approach to splitting the computations between multiple computing cores/units, which is a significant advantage considering the relatively large dimensionality of HARDI data.

4. Numerical solution

To simplify the solution of (15), we introduce two auxiliary variables u and v, and replace the original optimization problem by an equivalent, equality-constrained one. Specifically,

$$\begin{aligned} &\min_{f,u,v} \Big\{ \frac{1}{2}\|\Phi f - s\|_2^2 + \lambda\|u\|_1 + \mu\|v\|_a^2 + \nu\|v\|_{TV} + \varphi_{\geq}(u) \Big\}, \\ &\qquad\qquad \text{subject to } f = u,\ f = v \end{aligned} \tag{16}$$

Note that in (16) the minimization is carried out with respect to three variables, namely f, u, and v. A standard approach to solving such equality-constrained problems is based of the use of augmented Lagrangian methods [61]. Particularly, for the case at hand, this approach amounts to the following iterations (starting from some intitial values of Lagrange multipliers p_u^0 and p_v^0, e.g., $p_u^0 = p_v^0 = 0$).

$$\begin{aligned} (f^{(t+1)}, u^{(t+1)}, v^{(t+1)}) = \arg\min_{f,u,v} \Big\{ &\frac{1}{2}\|\Phi f - s\|_2^2 + \lambda\|u\|_1 + \mu\|v\|_a^2 + \nu\|v\|_{TV} + \\ &+\varphi_{\geq}(u) + \frac{\delta_u}{2}\|f - u + p_u^{(t)}\|_2^2 + \frac{\delta_v}{2}\|f - v + p_v^{(t)}\|_2^2 \Big\} \\ p_u^{(t+1)} = p_u^{(t)} + f^{(t+1)} - u^{(t+1)} \\ p_v^{(t+1)} = p_v^{(t)} + f^{(t+1)} - v^{(t+1)} \end{aligned} \tag{17}$$

where t stands for an iteration index, while $\delta_u > 0$ and $\delta_v > 0$ are some positive constants[3].

Finally, in ADMM, the concurrent minimization with respect to f, u, and v is replaced by sequential minimization with respect to f, u, and v independently. In this case, starting

[3]Note that the algorithm is guaranteed to converge for any positive δ_u and δ_v. In the present work we use $\delta_u = \delta_v = 0.5$.

from some $f^{(0)} = u^{(0)} = v^{(0)}$, the update proceeds according to:

(18)
$$\textbf{Step1}: f^{(t+1)} = \arg\min_{f} \Big\{ \frac{1}{2}\|\Phi f - s\|_2^2 + \frac{\delta_u}{2}\|f - u^{(t)} + p_u^{(t)}\|_2^2 + \frac{\delta_v}{2}\|f - v^{(t)} + p_v^{(t)}\|_2^2 \Big\}$$

(19)
$$\textbf{Step2}: u^{(t+1)} = \arg\min_{u} \Big\{ \frac{\delta_u}{2}\|f^{(t+1)} - u + p_u^{(t)}\|_2^2 + \lambda\|u\|_1 + \varphi_{\geq}(u) \Big\}$$

(20)
$$\textbf{Step3}: v^{(t+1)} = \arg\min_{v} \Big\{ \frac{\delta_v}{2}\|f^{(t+1)} - v + p_v^{(t)}\|_2^2 + \mu\|v\|_a^2 + \nu\|v\|_{TV} \Big\},$$

followed by the "dual" step of updating the Lagrange multipliers p_u and p_v, as specified in (17). Although the ADMM approach has effectively supplanted a single minimization problem (15) by three minimization problems, the latter admit much simpler solutions as detailed below.

4.1. **Solution to Step 1.** The optimization problem in (18) is a classical least-square (LS) problem which admits a closed-form solution as given by

$$f^{(t+1)} = \big(\Phi^T\Phi + (\delta_u + \delta_v)I_{J+1}\big)^{-1}\big(\Phi^T s + \delta_u(u^{(t)} - p_u^{(t)}) + \delta_v(v^{(t)} - p_v^{(t)})\big), \tag{21}$$

where I_{J+1} stands for a $(J+1)\times(J+1)$ identity matrix. To facilitate the computations, the inverse matrix $R \triangleq \big(\Phi^T\Phi + (\delta_u + \delta_v)I_{J+1}\big)^{-1}$ can be precomputed and stored before the reconstruction procedure is initiated.

4.2. **Solution to Step 2.** The optimization problem in (19) can be equivalently rewritten as

$$u^{(t+1)} = \arg\min_{u} \Big\{ \frac{1}{2}\|u - (f^{(t+1)} + p_u^{(t)})\|_2^2 + \frac{\lambda}{\delta_u}\|u\|_1 + \varphi_{\geq}(u) \Big\}. \tag{22}$$

In the absence of the indicator function $\varphi_{\geq}$, the solution to the above problem would be given by soft thresholding, *viz.* $u^{(t+1)} = \mathcal{S}_{\lambda/\delta_u}\{f^{(t+1)} + p_u^{(t)}\}$, with $\mathcal{S}_\tau(x) = \text{sign}(x)(|x| - \tau)_+$ (where the subscript + denotes the operation of keeping the positive part of the argument). As unexpected as it might sound, however, incorporating the indicator function actually simplifies the solution by replacing the operator $\mathcal{S}_\tau(x)$ with its positively rectified version $\mathcal{S}_\tau^+(x) = (x - \tau)_+$. Consequently, the solution to (19) is given by

$$u^{(t+1)} = \mathcal{S}^+_{\lambda/\delta_u}\{f^{(t+1)} + p_u^{(t)}\} = \big(f^{(t+1)} + p_u^{(t)} - \lambda/\delta_u\big)_+. \tag{23}$$

4.3. **Solution to Step 3.** To derive an update equation for $v^{(t)}$, it is convenient to rewrite the optimization problem in (20) in a slightly different form as given by

$$v^{(t+1)} = \arg\min_{v} \Big\{ \frac{1}{2}\|v - (f^{(t+1)} + p_v^{(t)})\|_2^2 + \frac{\mu}{\delta_v}\|v\|_a^2 + \frac{\nu}{\delta_v}\|v\|_{TV} \Big\} \tag{24}$$

Furthermore, since in our (simplified) notations $\|v\|_2^2 = \sum_{j=1}^{J+1} \|v_j\|_2^2$, then using the definitions of $\|v\|_a^2$ and $\|v\|_{TV}$ in (12) and (14), correspondingly, one can reexpress (24) in an equivalent form as

$$v^{(t+1)} = \arg\min_v \Big\{ \sum_{j=1}^{J+1} \Big(\frac{1}{2}\|v_j - (f^{(t+1)} + p_v^{(t)})_j\|_2^2 + \\ + \frac{\mu}{\delta_v}(1-\delta_{j,J+1})\|\mathbf{v}_j^T D v_j\|_2^2 + \frac{\nu}{\delta_v}\delta_{j,J+1}\|v_j\|_{TV} \Big) \Big\}, \tag{25}$$

where $(f^{(t+1)} + p_v^{(t)})_j$ denotes the j-th row of matrix $f^{(t+1)} + p_v^{(t)}$.

A closer look at (25) reveals that its cost functional consists of $J+1$ positive-valued terms, each of which depends on the variables v_j, $j = 1, 2, \ldots, J+1$, independently. Hence, the minimization over v can be replaced by minimizations of the $J+1$ summands in (25) with respect to their respective variables (i.e., v_j). In particular, for $j = 1, \ldots, J$, the resulting minimization problems are given by

$$v_j^{(t+1)} = \arg\min_{v_j} \Big\{ \frac{1}{2}\|v_j - (f^{(t+1)} + p_v^{(t)})_j\|_2^2 + \frac{\mu}{\delta_v}\|\mathbf{v}_j^T D v_j\|_2^2 \Big\}. \tag{26}$$

The above is a simple LS problem, which can be solved using spectral methods (i.e., by means of linear filtering) as detailed in the Appendix. Note that the cost function in (26) does not contain a TV term due to mutual exclusivity of the weights in definitions of $\|v\|_a^2$ and $\|v\|_{TV}$. For the same reason, optimization over v_{J+1} does not contain a fibre continuity term, resulting in

$$v_{J+1}^{(t+1)} = \arg\min_{v_{J+1}} \Big\{ \Big(\frac{1}{2}\|v_{J+1} - (f^{(t+1)} + p_v^{(t)})_{J+1}\|_2^2 + \frac{\nu}{\delta_v}\|v_{J+1}\|_{TV} \Big\}. \tag{27}$$

Note that the problem in (27) is a classical TV-regularized denoising problem [59], which can be efficiently solved by a variety of different algorithms. In the current work, we used the semi-implicit, fixed-point approach of [62] due to its impressive numerical stability and fast convergence.

5. Materials and Methods

5.1. **Sources of data.** The performance of the proposed and reference methods has been assessed using both computer-simulated and real-life data. Specifically, the simulated data were generated based on a standard Gaussian mixture model [23] with FA=0.8 and MD = $7 \cdot 10^{-4}$ mm^2/s (which corresponds to $\lambda_\parallel = 17 \cdot 10^{-4}$ mm^2/s and $\lambda_\perp = 3 \cdot 10^{-4}$ mm^2/s). The spatial dimensions of the data were set to be equal to $16 \times 16 \times 12$, while the directions of diffusion encoding were defined by the 2nd order tessellation of icosahedron restricted to the northern hemisphere of $\mathbb{S}^2$ (thus resulting in $K = 81$ sampling points). The data were designed so as to emulate a crossing of two cylindrically symmetric "fibres" of 8 voxels in diameter, with the crossing angle α in the range $\alpha \in [30^\circ, 90^\circ]$ with step size 5°. Subplots A and D of Fig. 2 depict the *theoretical* ODFs corresponding to the middle "layer" of two signal arrays which have been synthesized for $\alpha = 60^\circ$ and $\alpha = 45^\circ$, respectively.

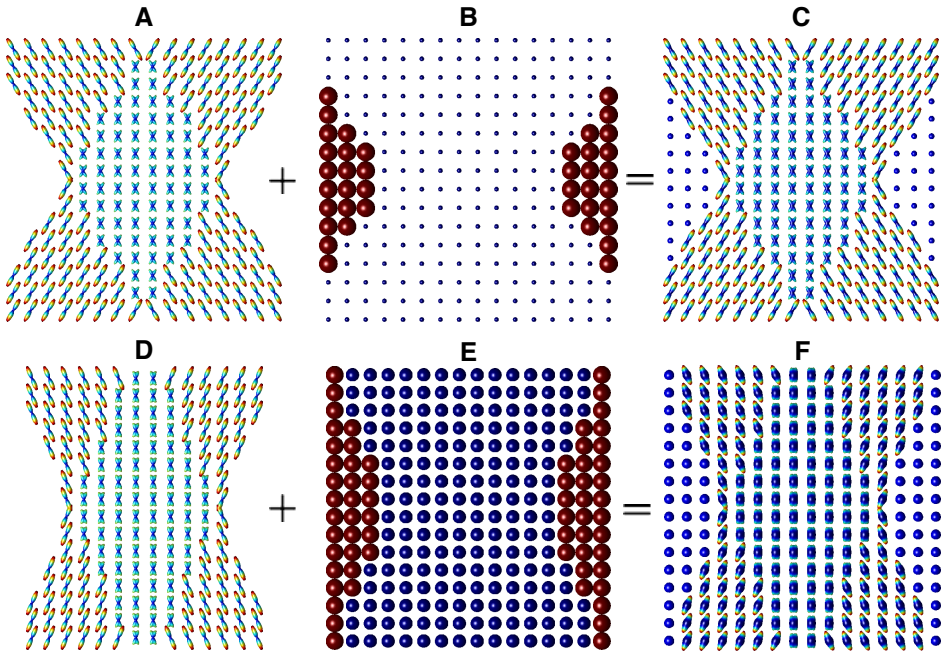

FIGURE 2. (Subplots A-C, left-to-right) The anisotropic, isotropic, and the combined field of ODFs associated with *in silico* HARDI signals generated with $b = 3000$ s/mm^2, $\alpha = 60°$, and $p_{iso} = 0.25$; (Subplots D-F) Same as above, only with $\alpha = 45°$ and $p_{iso} = 0.75$.

Additionally, to allow investigation of the effect of isotropic diffusion on the accuracy of SD-based reconstruction, the simulated signals have also been combined with a constant (i.e., isotropic) component of a variable magnitude. The diffusivity of the isotropic component was set to be equal to $8 \cdot 10^{-4}$ mm^2/s, whereas its partial volume fractions outside and inside of the "fibres" were set to 1 and p_{iso}, respectively, with $p_{iso} \in \{0, 0.25, 0.5, 0.75\}$. Two examples of the isotropic component are shown in Subplots B and E of Fig. 2 for the case of $p_{iso} = 0.25$ and $p_{iso} = 0.75$, correspondingly. (Note that, in the above subplots, the *absolute* size of the glyphs has been optimized for visualization, and thus it does not represent the actual values of the isotropic ODFs.) Finally, Subplot C of Fig. 2 depicts combined ODFs obtained as a result of the summation of the ODFs shown in Subplots A and B, while the result of the summation of the ODFs in Subplots D and E is displayed in Subplot F of the same figure.

The data simulation was repeated for two different values of b, *viz.* $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, with the corresponding SNRs of Rician noise[4] being 20 and 7. It should be noted that the above choice of b-values and their related SNRs is by no means arbitrary,

[4]In the case at hand, SNR was defined to be a ratio of the *mean* amplitude of a noise-free HARDI signal to the standard deviation of the Gaussian noise contaminating the complex MR readout.

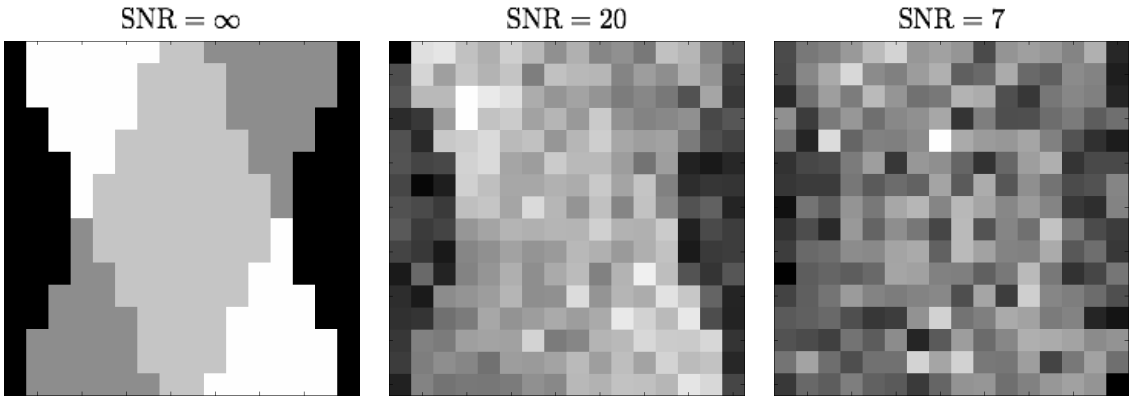

FIGURE 3. Example of a noise-free "axial slice" of the simulated HARDI data (left) along with its noise contaminated versions for SNR = 20 (middle) and SNR =7 (right).

but intended to imitate a real-life situation, in which increasing the value of b comes at the price of a substantially reduced SNR. (The effect of Rician noise is exemplified in Fig. 3 where a noise-free "axial slice" of the simulated HARDI signal (left) is shown along with its noise contaminated versions for SNR = 20 (middle) and SNR=7 (right).) In such a case, the adverse effect of measurement noises could effectively counterpoise the gain in signal bandwidth which comes with exploiting higher b-values. In the context of SD, it is therefore important to understand the trade-off between the values of b and SNR under various regularization schemes.

Additionally, real-life HARDI data were acquired by means of a 3T Siemens TrioTim MRI scanner. The diffusion-encoded scans were collected from a 45 y.o. healthy volunteer at 2 mm isotropic resolution with: $b \in \{1000, 3000\}$ s/mm^2, TR=6300 ms, and TE = 85 ms. The transversal dimensions of the scans were equal to 128×128 pixels, with a total of 50 axial slices used. The diffusion-encoding directions $\{\mathbf{u}_k\}_{k=1}^K$ (with $K = 64$) were defined using the method of generalized spirals [63], which provides a closed-form solution to the problem of quasi-uniform sampling of $\mathbb{S}^2$.

5.2. **Reference methods.** In this work, we compared the performance of the proposed method against that of a number of alternative approaches. The first of these approaches was the positively constrained LS algorithm of [36], which will be referred below to as CSD. Note that this method disregards the presence of isotropic diffusion, and its solutions can be produced by solving (15) upon replacing Φ by H and setting $\lambda = \mu = \nu = 0$. In the case when $\lambda > 0$, CSD transforms into the method of [42], which was the second reference method (referred below to as Min-$\mathcal{L}_1$) used in our comparative study. Note that, just like CSD, Min-$\mathcal{L}_1$ disregards *both* the presence of isotropic diffusion and the spatial-domain regularity of estimated fODFs. While still ignoring the effect of isotropic diffusion, the method of [52] takes advantage of an original way to constrain the spatial behaviour of

reconstructed fODFs using a fibre continuity model. This method (which we refer to below as CSD-FC) can be also described by the minimization problem in (15) upon replacing Φ by H and setting $\lambda = \nu = 0$. Further, setting $\mu = 0$ in (15) would result in the SD algorithm recently proposed in [64]. This method explicitly accounts for the presence of isotropic diffusion, while constraining the spatial regularity of its related IDM. However, the method of [64] (referred below to as Min-TV-$\mathcal{L}_1$) does not enforce the spatial continuity of fibre orientations, as it is done in the case of the proposed algorithm. In [44], the authors proposed an SD approach based on the famed Lucy-Richardson deconvolution procedure (referred to as dRL). Although disregarding the spatial-domain regularity of the estimated fODFs, dRL had been particularly designed to suppress the influence of isotropic diffusion on the deconvolution results. Finally, for the sake of notational convenience, the SD method proposed in this paper will be referred below to as SCSD (standing for spatially constrained sparse deconvolution).

5.3. **Comparison metrics.** To quantitatively assess the results produced by different SD routines, a total of four comparison metrics have been employed. The first of these metrics is *average angular error* (AAE) defined as

$$\mathrm{AAE} = \frac{180^\circ}{\pi}\mathcal{E}\{\arccos(\mathbf{v}_0 \cdot \tilde{\mathbf{v}})\}, \tag{28}$$

where $\mathbf{v}_0$ and $\tilde{\mathbf{v}}$ stand for an original fibre orientation and its estimate respectively and $\mathcal{E}$ denotes the operator of statistical expectation (which is commonly approximated by a sample mean in practice). It should be noted that a standard way to infer the value(s) of $\tilde{\mathbf{v}}$ from an estimated fODF is through finding the local maxima of the latter. Alternatively, one could use the estimated fODF to fit to it a model of the form $\sum_{l=1}^{M} \exp(\mathbf{u}^T A_l \mathbf{u})$ (with M denoting the number of fibres and A_l being a symmetric 3×3 matrix) and then associate $\tilde{\mathbf{v}}$ with the direction of the major eigenvector of A_l. When the number of fibres M is known as *a priori* (as it is the case with computer-simulated studies), the second approach tends to produce more accurate results. For this reason, it was exploited in this paper.

Additional metrics used in our comparative analysis were *true positive* (TP) and *false positive* (FP) rates of detection of the correct number of fibres at each voxel. Specifically, TP can be defined to be a percentage of data voxels, in which the number of local maxima of an estimated fODF coincided with the true number of simulated fibres. To render the computation of TP more robust towards overestimation errors, at each voxel the local maxima have been subjected to hard thresholding, set at the level of 20% of their peak value. The FP metric, on the other hand, was particularly designed to assess the extent of overestimation of the true number of fibres [52, 44]. Formally, FP can be defined as

$$\mathrm{FP} = \mathcal{E}\{(\tilde{M} - M_0)_+\}, \tag{29}$$

where M_0 and $\tilde{M}$ denote the true number of fibres and the number of local maxima in reconstructed fODFs. Note that, similarly to the case of TP, before evaluating FP, all the local maxima had been subjected to the same hard thresholding procedure.

Finally, one of the principal outputs of the proposed method is an estimate of the IDM, which is a scalar-valued function of the spatial coordinate. In the case at hand, the quality

of reconstructed IDMs can be assessed qualitatively based on the notion of image contrast. The latter can be defined as follows. First, we partition the entire image domain Ω into two subdomains, Ω_{in} and Ω_{out} (with $\Omega = \Omega_{in} \cup \Omega_{out}$), which encompass the regions occupied by the simulated fibres and purely isotropic diffusion (i.e., no fibres present), respectively. Also, we denote by $f_{iso}(\Omega_{in})$ and $f_{iso}(\Omega_{out})$ the restrictions of f_{iso} to the corresponding subdomains. Then, letting μ_{in} and μ_{out} (resp. σ_{in} and σ_{out}) denote the mean values (resp. the standard deviations) of $f_{iso}(\Omega_{in})$ and $f_{iso}(\Omega_{out})$, correspondingly, a measure of contrast C can be defined as given by

$$C = 2\,\frac{|\mu_{in} - \mu_{out}|}{\sigma_{in} + \sigma_{out}}. \tag{30}$$

It is worthwhile noting that, the (real) IDMs used in the experimental study are piecewise constant "images", in which case $\sigma_{in} = \sigma_{out} = 0$, resulting in $C = \infty$. Thus, in the case of estimated IDMs, it is reasonable to assume that higher values of C represent more accurate estimates of the original IDMs.

As a closing remark we note that, among all the SD methods under comparison, only Min-TV-$\mathcal{L}_1$ and SCSD are capable (by design) of computing the estimates of IDM, whereas neither of the remaining methods is endowed with the same capability. However, both CSD and Min-$\mathcal{L}_1$, as well as CSD-FC, have been proposed as means for estimation of the fODF f_a, in which case the remainder $H\tilde{f}_a - s$ (with $\tilde{f}_a$ being an estimate of f_a) is likely to represent estimation errors *along with* a contribution from f_{iso}. Furthermore, assuming the estimation errors to have a zero mean value, a reasonable estimate $\tilde{f}_{iso}$ of f_{iso} can be obtained by averaging the remainder $H\tilde{f}_a - s$ over the spherical coordinate. Formally, one can define $\tilde{f}_{iso}$ to be

$$\tilde{f}_{iso} \simeq \frac{A}{K}\,\mathbf{1}^T\big(H\tilde{f}_a - s\big), \tag{31}$$

where $\mathbf{1}$ denotes a K-dimensional (column) vector of ones and $A > 0$ is a proportionality constant that could be set to 1, since it does not have any effect on the value of C in (30). We also note that, by its design, the dRL algorithm of [44] aims to recover an FOF (rather than an fODF) as given by (5), thereby making the estimation of IDMs according to (31) inadequate. Consequently, quantitative comparisons in terms of C have been performed for all the methods under consideration, except for dRL.

6. Results

6.1. **Computer simulations.** The left and right columns of subplots in Fig. 4 show the values of AAE as a function of crossing angle α, which have been obtained with the proposed and reference methods for $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively. The regularization parameters for Min-$\mathcal{L}_1$, CSD-FC, Min-TV-$\mathcal{L}_1$, and SCSD were set to optimize the overall performance of the algorithms, namely: 1) $\lambda = 0.01$ for Min-$\mathcal{L}_1$, 2) $\mu = 0.01$ for CSD-FC, 3) $\lambda = 0.07$, $\nu = 0.01$ for Min-TV-$\mathcal{L}_1$, and 4) $\lambda = 0.03$, $\nu = 0.01$, $\mu = 0.4$ for SCSD. The dRL method has been reproduced following its description in [44]. One can see that all the error curves, with the exception of that of dRL, exhibit the expected behaviour

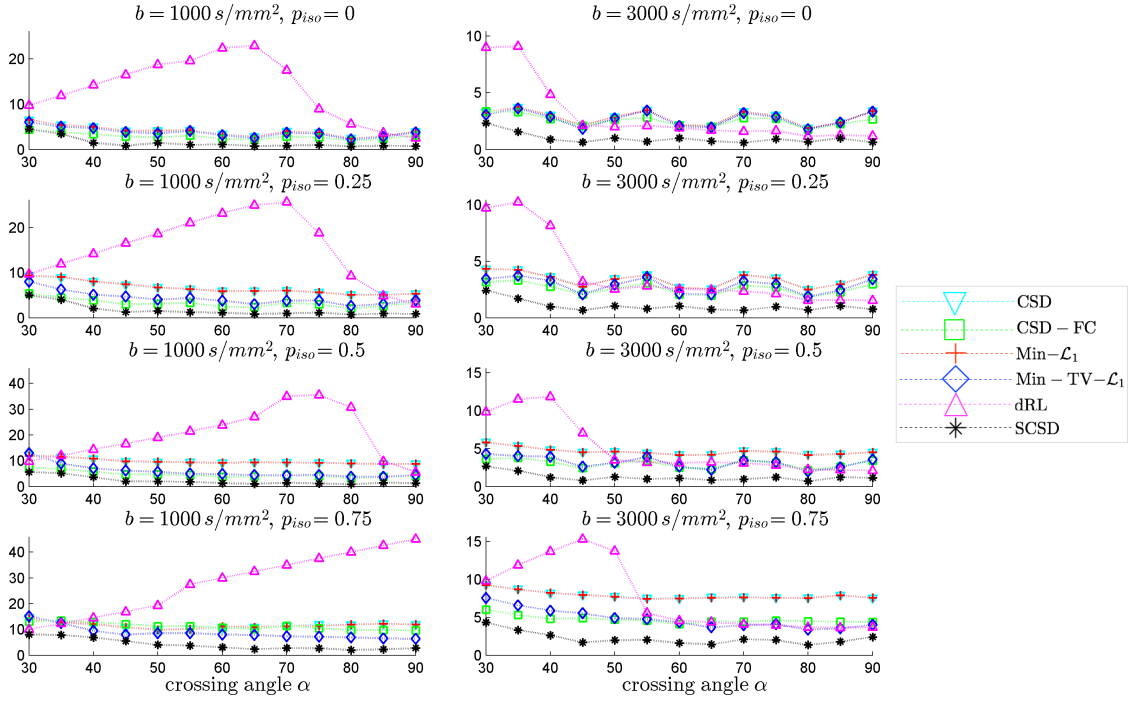
$b = 1000\, s/mm^2$, $p_{iso} = 0$
$b = 3000\, s/mm^2$, $p_{iso} = 0$
$b = 1000\, s/mm^2$, $p_{iso} = 0.25$
$b = 3000\, s/mm^2$, $p_{iso} = 0.25$
$b = 1000\, s/mm^2$, $p_{iso} = 0.5$
$b = 3000\, s/mm^2$, $p_{iso} = 0.5$
$b = 1000\, s/mm^2$, $p_{iso} = 0.75$
$b = 3000\, s/mm^2$, $p_{iso} = 0.75$
crossing angle α
crossing angle α
CSD
CSD − FC
Min−$\mathcal{L}_1$
Min − TV−$\mathcal{L}_1$
dRL
SCSD

FIGURE 4. (Left columns of subplots) AAE produced by the tested methods for different α, p_{iso}, and $b = 1000$ s/mm^2; (Right columns of subplots) AAE produced by the tested methods for different α, p_{iso}, and $b = 3000$ s/mm^2.

where AAE decreases with an increase in α. Moreover, despite the substantially worse noise conditions for $b = 3000$ s/mm^2, all the tested methods (again, with the exception of dRL) demonstrate better performance for $b = 3000$ s/mm^2, as compared to the case of $b = 1000$ s/mm^2. In all the cases, however, the proposed SCSD method shows considerably better performance in comparison to the alternative solutions, with the "second best" results produced by CSD-FC for lower values of p_{iso} and by Min-TV-$\mathcal{L}_1$ for higher values of p_{iso}.

Before proceeding any further, one additional comment is in order regarding the behaviour of the AAE curves obtained with dRL. Specifically, one can see that, for $b = 1000$ s/mm^2, the AAE is minimized for smaller values of α, which is rather a counter-intuitive result. To understand why this happens, it is instructive to examine the behaviour of the TP curves produced by dRL (see Fig. 5). Specifically, one can see that for $b = 1000$ s/mm^2 and $p_{iso} = 0$, dRL is incapable of resolving the crossing fibres of the numerical phantom for $\alpha \leq 60°$, with the resolvability problem becoming progressively worse with an increase in p_{iso}. In this case, the values of AAE effectively "mirror" those of α up to the point when dRL starts detecting the correct number of fibres, after which AAE becomes a decreasing function of the fibre crossing angle. It also deserves noting that for $b = 3000$ s/mm^2, dRL

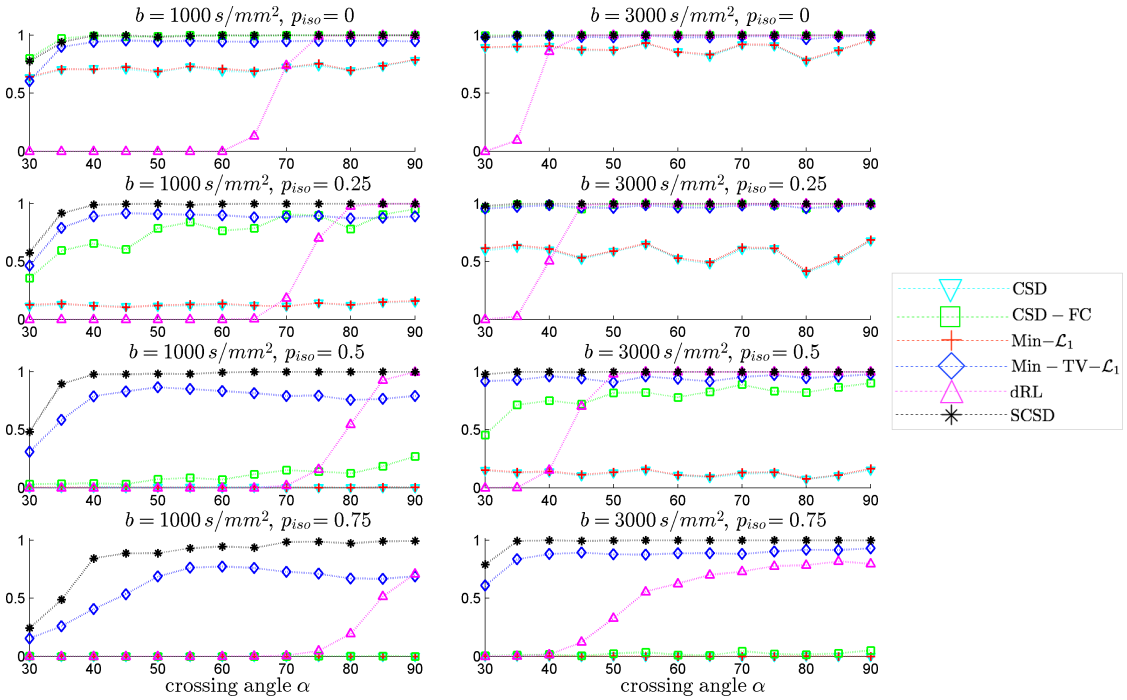
$b = 1000\, s/mm^2$, $p_{iso} = 0$
$b = 3000\, s/mm^2$, $p_{iso} = 0$
$b = 1000\, s/mm^2$, $p_{iso} = 0.25$
$b = 3000\, s/mm^2$, $p_{iso} = 0.25$
$b = 1000\, s/mm^2$, $p_{iso} = 0.5$
$b = 3000\, s/mm^2$, $p_{iso} = 0.5$
$b = 1000\, s/mm^2$, $p_{iso} = 0.75$
$b = 3000\, s/mm^2$, $p_{iso} = 0.75$
crossing angle α
crossing angle α
CSD
CSD − FC
Min−$\mathcal{L}_1$
Min − TV−$\mathcal{L}_1$
dRL
SCSD

FIGURE 5. (Left columns of subplots) TP produced by the tested methods for different α, p_{iso}, and $b = 1000$ s/mm^2; (Right columns of subplots) TP produced by the tested methods for different α, p_{iso}, and $b = 3000$ s/mm^2.

demonstrates considerably improved performance in terms of AAE, even surpassing CSD and Min-$\mathcal{L}_1$ for $\alpha \geq 50^\circ$ and $p_{iso} \geq 0.5$. However, in all the alternative scenarios, the performance of dRL remains inferior to that of other methods under comparison. A possible explanation to this fact could be that the Lucy-Richardson algorithm exploited by dRL aims at recovering a maximum likelihood estimate under the assumption on measurement noise to be Poissonian. However, such a noise model can hardly be a good approximation to Rician distribution, which is inherent in MRI.

Fig. 5 shows the values of TP which have been obtained using the proposed and reference methods for $b = 1000$ s/mm^2 (left column of subplots) and $b = 3000$ s/mm^2 (right column of subplots). One can see that, in the absence of isotropic diffusion (i.e., for $p_{iso} = 0$), the proposed SCSD method performs comparably to CSD-FC for both values of b (with Min-TV-$\mathcal{L}_1$ being the next "best performer"). Yet, the moment p_{iso} starts growing above zero, not only CSD and Min-$\mathcal{L}_1$, but also CSD-FC start loosing their ability to accurately estimate the true number of simulated fibres. Note that this is an expected result, considering the fact that neither of the above methods is endowed with facilities to explicitly account for the presence of isotropic diffision, in which case over-estimation errors become

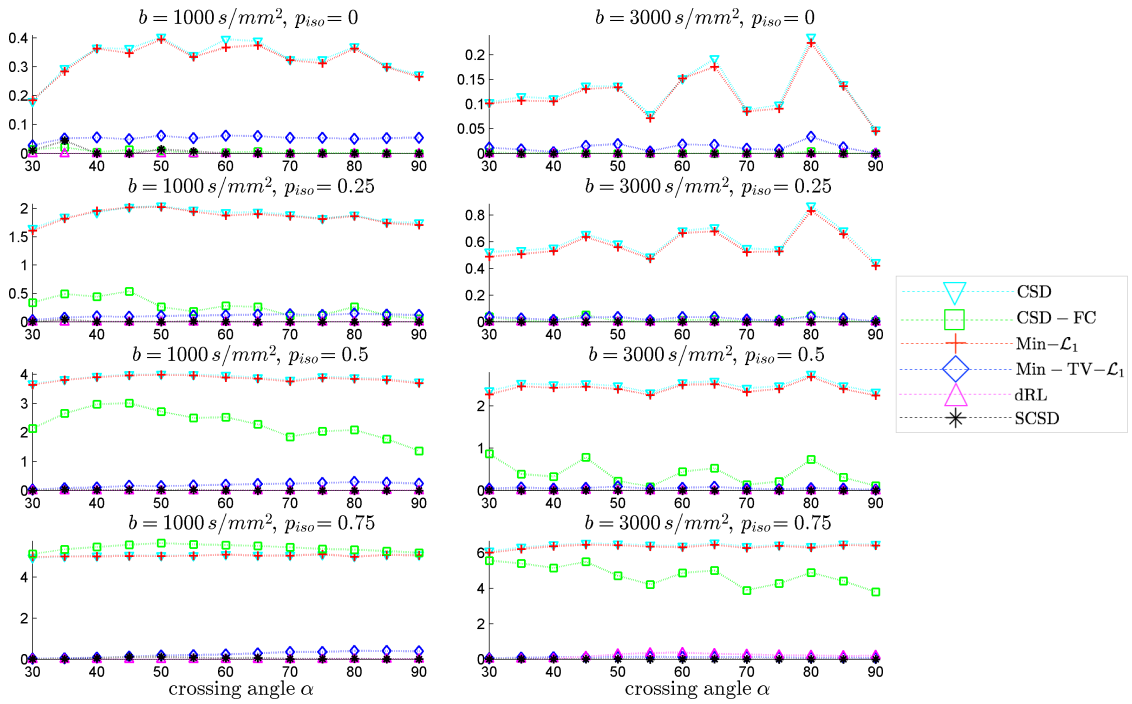
$b = 1000\, s/mm^2$, $p_{iso} = 0$
$b = 3000\, s/mm^2$, $p_{iso} = 0$
$b = 1000\, s/mm^2$, $p_{iso} = 0.25$
$b = 3000\, s/mm^2$, $p_{iso} = 0.25$
$b = 1000\, s/mm^2$, $p_{iso} = 0.5$
$b = 3000\, s/mm^2$, $p_{iso} = 0.5$
$b = 1000\, s/mm^2$, $p_{iso} = 0.75$
$b = 3000\, s/mm^2$, $p_{iso} = 0.75$
crossing angle α
crossing angle α
CSD
CSD − FC
Min−$\mathcal{L}_1$
Min − TV−$\mathcal{L}_1$
dRL
SCSD

FIGURE 6. (Left columns of subplots) FP produced by the tested methods for different α, p_{iso}, and $b = 1000$ s/mm^2; (Right columns of subplots) FP produced by the tested methods for different α, p_{iso}, and $b = 3000$ s/mm^2.

inevitable (see below). At the same time, for $b = 3000$ s/mm^2, the proposed SCSD algorithm provides an ideal detection rate of one for $p_{iso} \in \{0, 0.25, 0.5\}$ and $\alpha \geq 30^\circ$, as well as for $p_{iso} = 0.75$ and $\alpha \geq 35^\circ$. It deserves noting that, owing to its properly accounting for the effect of isotropic diffusion, the performance of Min-TV-$\mathcal{L}_1$ is only marginally inferior to that of SCSD for $b = 3000$ s/mm^2 and $p_{iso} \leq 0.5$. One can also see that a nearly ideal TP rate is reached by the dRL algorithm as well. However, it does not happen until after α reaches relatively large values (e.g., for $\alpha \geq 50^\circ$ with $p_{iso} = 0.5$ and $b = 3000$ s/mm^2).

One of the principle applications of HARDI is in multi-fibre tractography, in which case SD can be used to improve the resolvability of multiple fibre tracts within each given voxel in a region of interest (ROI). In this situation, overestimating the number of fibres is likely to produce spurious fibre tracts, thereby rendering the resulting reconstructions unreliable. Unfortunately, unless properly regularized, some SD routines tend to amplify the effect of noise, which in turn results in numerous false (local) maxima in reconstructed fODFs – the maxima that can be easily confused with the true modes of the latter. For this reason, it is important to compare the performance of SD routines in terms of the FP metric. Such comparative results are summarized in Fig. 6, the left and right columns of which correspond to the cases of $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively. One can

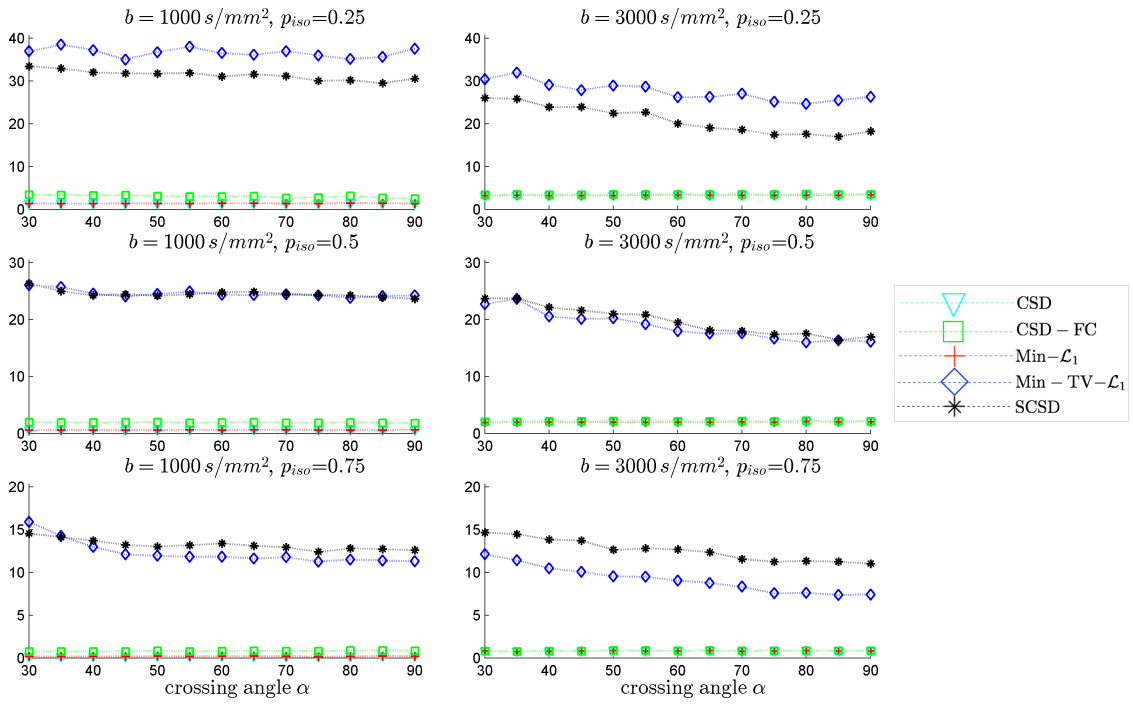

FIGURE 7. (Left columns of subplots) C produced by the tested methods for different α, p_{iso}, and $b = 1000$ s/mm^2; (Right columns of subplots) C produced by the tested methods for different α, p_{iso}, and $b = 3000$ s/mm^2.

see that, in this case, the worst results are produced by CSD and Min-$\mathcal{L}_1$ for all values of b and α. The CSD-FC algorithm, on the other hand, provides a close to zero FP rate for $b = 3000$ s/mm^2 and $p_{iso} \in \{0, 0.25\}$. Unfortunately, its performance deteriorates for higher values of p_{iso} (which is particularly noticeable for $b = 1000$ s/mm^2). Surprisingly, dRL seems to provide an ideal FP rate of zero for all values of b and α under consideration. However, from our analysis of Fig. 5 it is not hard to see that it happens only because this method tends to underestimate the true number of simulated fibres. Finally, one can also see that, for $b = 3000$ s/mm^2, Min-TV-$\mathcal{L}_1$ provides a fairly small FP rate, whereas the proposed SCSD algorithm succeeds to attain the ideal FP rate of zero for all values of p_{iso} and α.

The results of our final quantitative comparison are summarized in Fig. 7, which shows the values of contrast C obtained using different SD methods under comparison for $b = 1000$ s/mm^2 (left column of subplots) and $b = 3000$ s/mm^2 (right column of subplots). Predictably enough, the best contrast is achieved by the Min-TV-$\mathcal{L}_1$ and SCSD algorithms, owing to their inherent ability to account for the presence of isotropic diffusion. Moreover, out of the two, the proposed SCSD algorithm yields the higher values of C for all simulated scenarios. An additional illustration of the effect of incorporation and spatial regularization

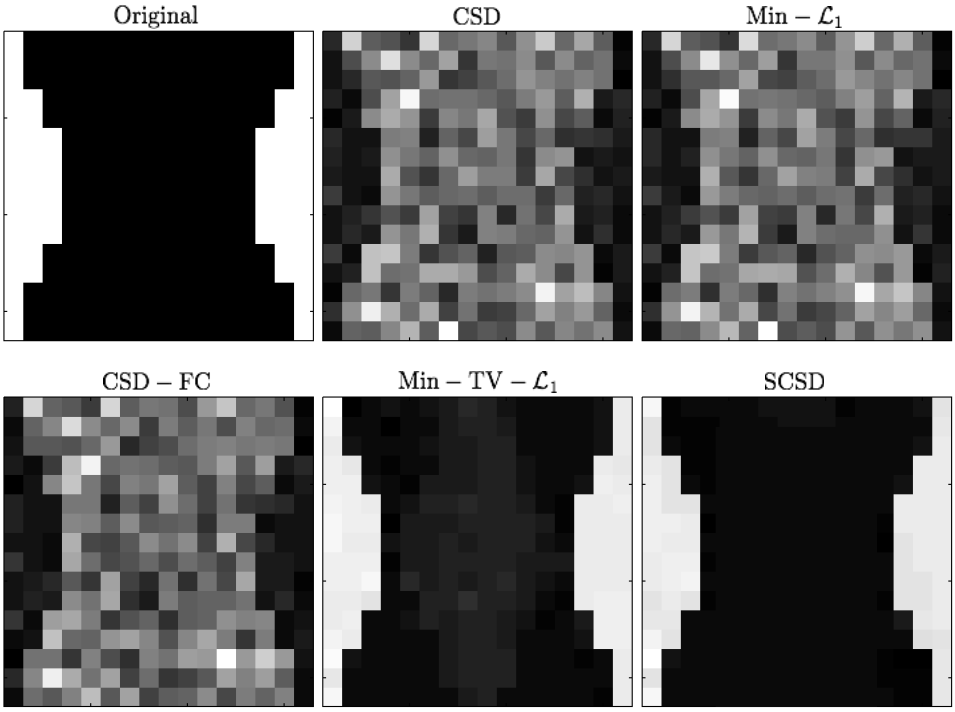

FIGURE 8. An "axial" slice of the IDMs recovered by different SD methods under comparison for $b = 3000$ s/mm^2, $\alpha = 45^\circ$, and $p_{iso} = 0.5$.

of the isotropic diffusion component is provided in Fig. 8, which depicts a 2-D "axial" slice of the IDMs reconstructed by different SD methods under comparison for $b = 3000$ s/mm^2, $\alpha = 45^\circ$, and $p_{iso} = 0.5$. (Note that, for the sake of the clarity of visualization, the IDMs in Fig. 8 have been normalized so as to make their minimum and maximum values correspond to black and white pixel values, respectively.) One can see that the IDM reconstruction produced by SCSD is virtually indistinguishable from the original IDM (as shown in the upper, leftmost subplot of the figure), with the second best result produced by the Min-TV-$\mathcal{L}_1$ algorithm. At the same time, neither CSD, Min-$\mathcal{L}_1$ nor CSD-FC can attain a comparable accuracy of estimation of the spatial pattern of isotropic diffusion, as represented by the original IDM. In particular, even though their respective reconstructions do bear some global resemblance to the original IDM, the level of estimation errors is too high to deem these reconstructions useful.

6.2. **In vivo data experiments.** As the next step of our experimental study, real-life estimation has been performed using *in vivo* diffusion data (see Section 5.1). For the sake of the clarity of visualization, we restrict demonstration of the obtained reconstructions to 2-D views (aka "slices"), two examples of which are depicted in Fig 9. In particular,

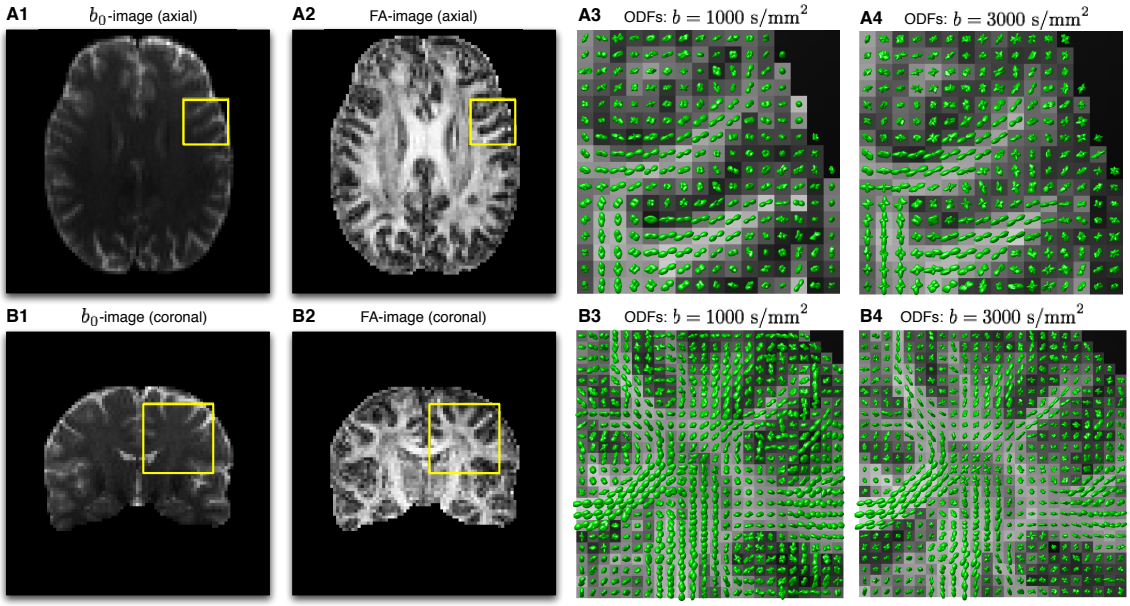

FIGURE 9. (Subplots A1-A2) b_0-image and FA images of an axial slice of an *in vivo* scan used for presentation of experimental results; (Subplots A3-A4) ODFs corresponding to the localized region indicated by the yellow squares in Subplots A1-A2 for $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively; (Subplots B1-B2) b_0-image and FA images of a coronal slice from the same scan; (Subplots B3-B4) ODFs corresponding to the localized region indicated by the yellow squares in Subplots B1-B2 for $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively.

Subplots A1-A2 of the figure display an axial slice of the acquired b_0-volume and its corresponding fractional anisotropy (FA) image, respectively, while Subplots B1-B2 show a coronal slice from the same 3-D volume along with its associated FA image, in the same order. Additionally, Subplots A3-A4 of Fig. 9 show the ODFs corresponding to the localized regions indicated by the yellow squares in Subplots A1-A2 for the case of $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively. (To facilitate the analysis, the ODFs are shown over the background of their associated FA values.) Analogous results pertaining to the coronal view are shown in Subplots B3-B4 of the figure. It should be noted that, in both cases, the ODFs have been computed by means of the FRACT algorithm of [65], which seems to provide a reasonable balance between the robustness of more traditional QBI [25] and the high resolution gain of its solid-angle formulation [27].

Fig. 10 depicts the reconstructions of fODFs obtained using the CSD (Subplot A), Min-$\mathcal{L}_1$ (Subplot B), CSD-FC (Subplot C), dRL (Subplot D), Min-TV-$\mathcal{L}_1$ (Subplot E), and

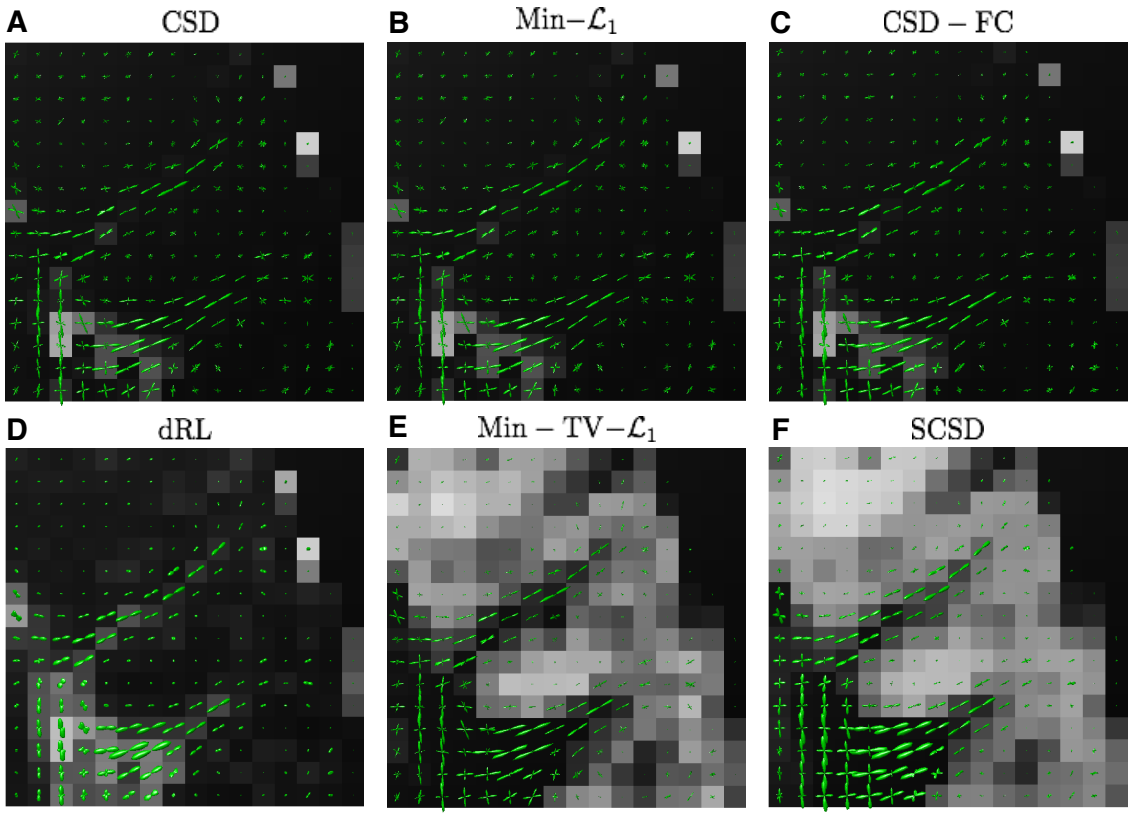

FIGURE 10. Estimated fODFs computed by CSD (A), Min-$\mathcal{L}_1$ (B), CSD-FC (C), dRL (D), Min-TV-$\mathcal{L}_1$ (E), and SCSD (F) for the axial view and $b = 1000$ s/mm^2. The fODFs are superimposed over the background of their related IDMs.

SCSD (Subplot F) for the axial view and $b = 1000$ s/mm^2. As opposed to Subplots A3-4 and B3-4 in Fig. 9, the fODFs in Fig. 10 are shown superimposed over the values of their corresponding IDMs. Analysing these results reveals the principal drawback of SD methods which disregard the effect of isotropic diffusion. In particular, not only CSD and Min-$\mathcal{L}_1$ but also CSD-FC tend to yield spurious estimates of fODFs in anatomical regions corresponding to the cortical grey matter – the result which stands at odds with the fact that these regions are known to be devoid of neural fibre bundles. Further, although being capable of coping with the presence of isotropic diffusion, the dRL algorithm is "blending" f_a and f_{iso}, which effectively impairs the angular resolution, and therefore the resolvability of crossing fibre tracts. Moreover, neither of the aforementioned methods has been found to be capable of reliably recovering the IDMs. At the same time, both Min-TV-$\mathcal{L}_1$ and SCSD yield anatomically consistent reconstructions of the IDMs, with much less noisy

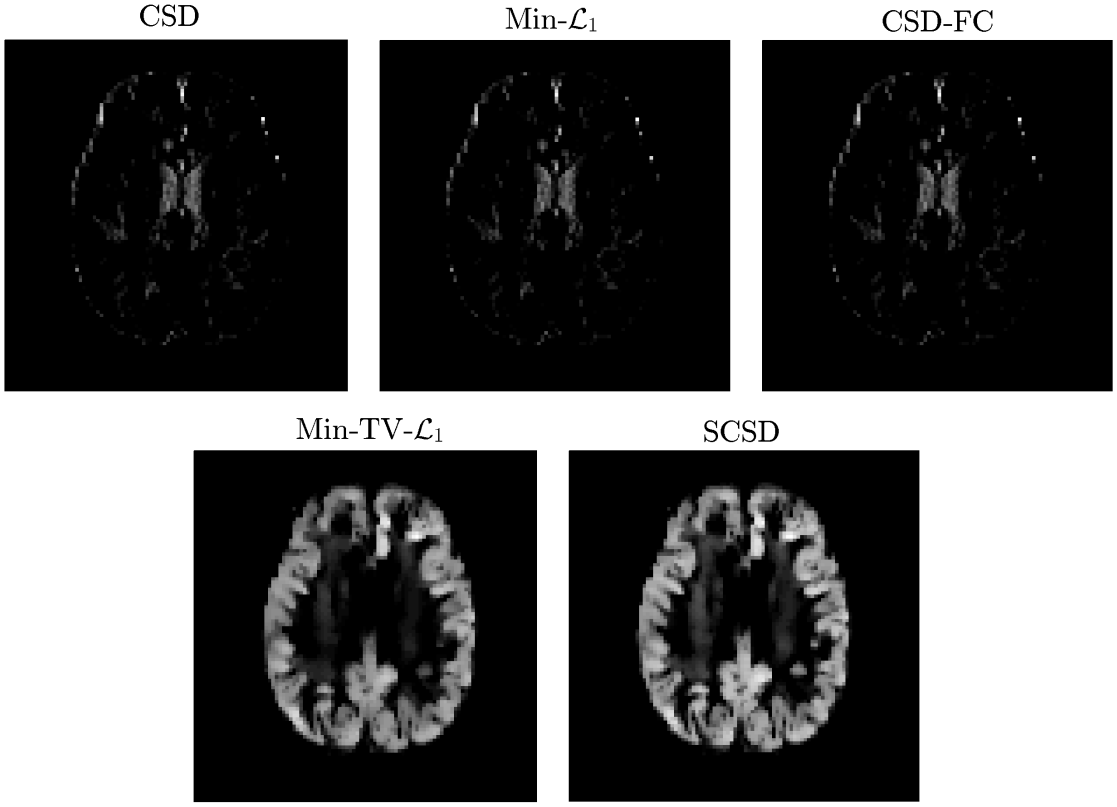

FIGURE 11. "Axial" IDMs estimated by various SD methods for $b = 1000$ s/mm^2.

results obtained in the case of SCSD. (This point is further illustrated by Fig. 11 which shows the "zoomed-out" IDMs recovered by the SD methods under comparison for the case of $b = 1000$ s/mm^2.) Moreover, a closer inspection of Fig. 10 reveals that the fODFs estimated by means of SCSD are characterized by a smoother and more consistent spatial variability (owing to the fibre continuity constraint), thereby exhibiting a better adherence to the expected connectivity within an *in vivo* brain.

The reconstructions obtained for the same axial view and $b = 3000$ s/mm^2 are depicted in Fig. 12, whose composition is identical to that of Fig. 10. Although fairly close in appearance to the previous case, these reconstructions allow us to make a number of important observations. First of all, as compared to the case of $b = 1000$ s/mm^2, a wider bandwidth of the HARDI signals at $b = 3000$ s/mm^2 leads to a better angular resolution, which is particularly noticeable in the case of dRL. Moreover, despite considerably worse noise conditions, the fODF reconstructions obtained by means of Min-TV-$\mathcal{L}_1$ and SCSD have much less residual noise over the areas occupied by cortical grey matter, where isotropic diffusion is expected to prevail. This fact indicates the effectiveness of the regularization schemes

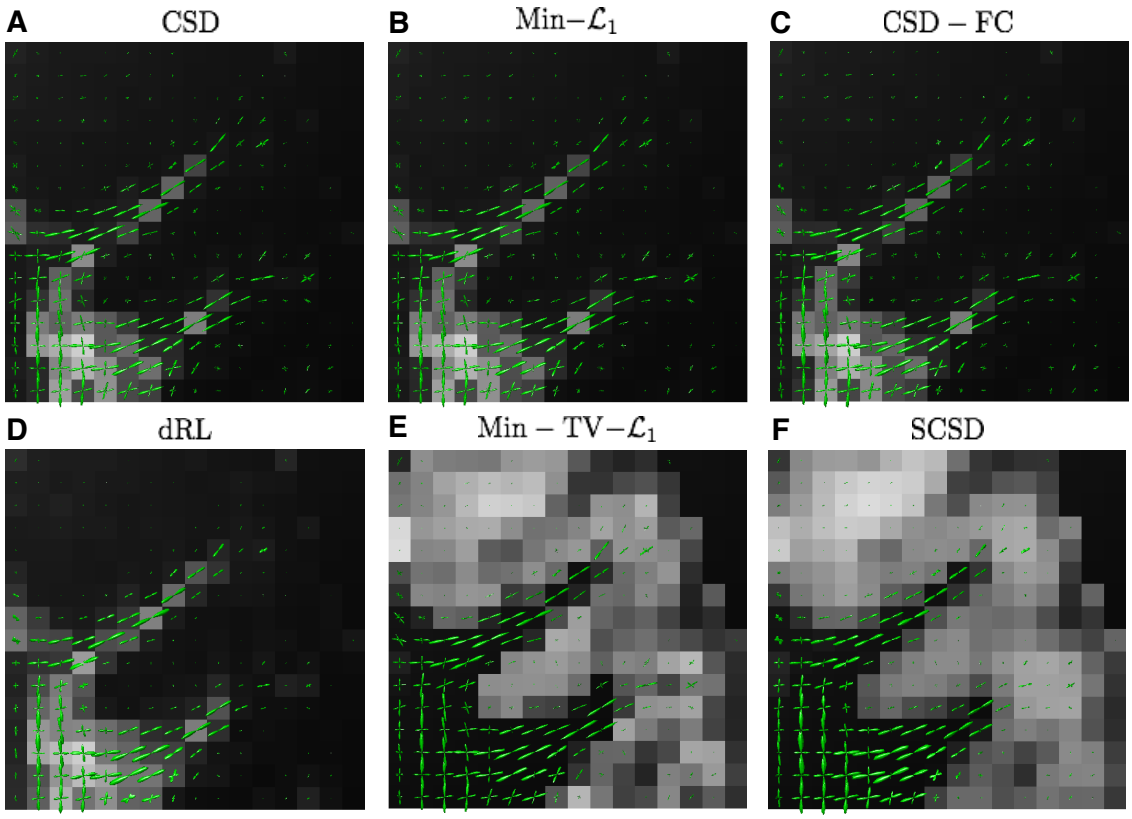

FIGURE 12. Estimated fODFs computed by CSD (A), Min-$\mathcal{L}_1$ (B), CSD-FC (C), dRL (D), Min-TV-$\mathcal{L}_1$ (E), and SCSD (F) for the axial view and $b = 3000$ s/mm^2. The fODFs are superimposed over the background of their related IDMs.

exploited by these SD methods. Finally, a closer inspection of the glyphs in Fig. 12 reveals that the fODF reconstructions yielded by SCSD demonstrate a better spatial smoothness and anatomical consistency. The same observations can be made in the case of the coronal view, as shown in Figs. 13 and 14 for $b = 1000$ s/mm^2 and $b = 3000$ s/mm^2, respectively. (Note that the local region represented by these figures has been chosen according to the results in [44], which also provides indication of specific fibre bundles within the selected regions of interest).

7. Discussion and conclusions

In the present paper, we introduced a novel approach to the problem of non-blind SD of HARDI data. As opposed to many alternative methods of SD, the proposed algorithm can explicitly account for the effect of isotropic diffusion, which makes it capable of performing

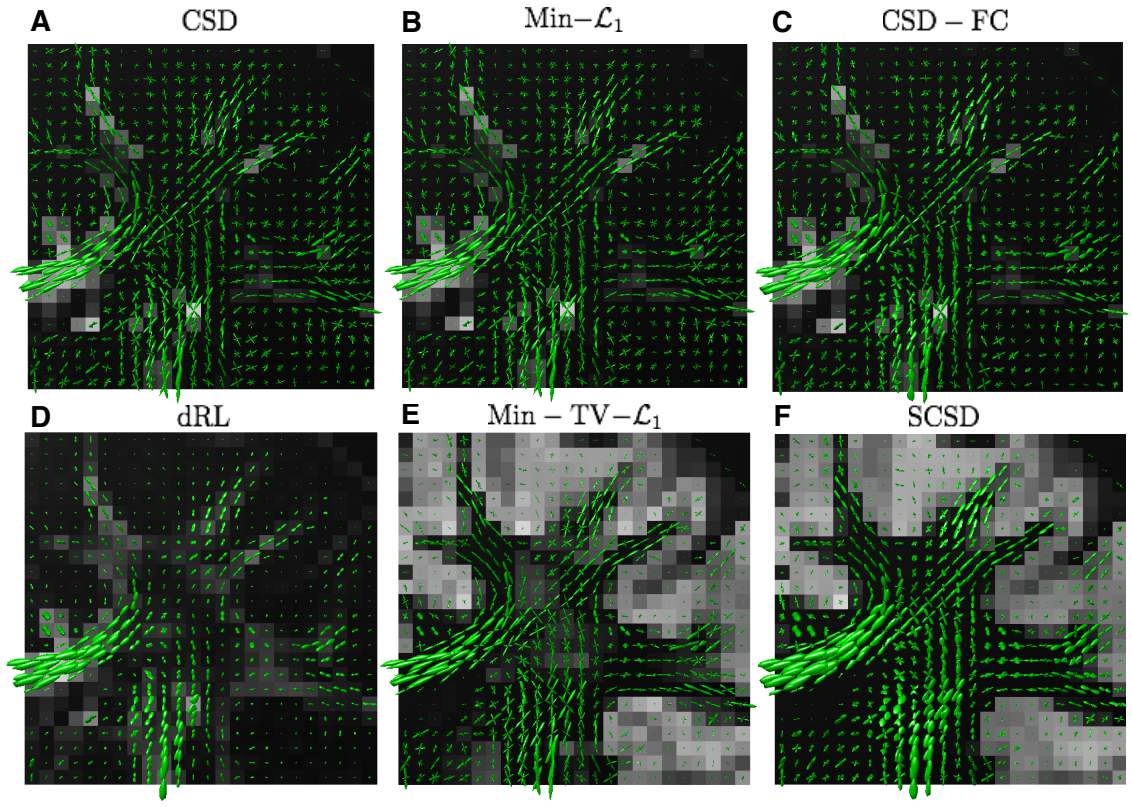

FIGURE 13. Estimated fODFs computed by CSD (A), Min-$\mathcal{L}_1$ (B), CSD-FC (C), dRL (D), Min-TV-$\mathcal{L}_1$ (E), and SCSD (F) for the coronal view and $b = 1000$ s/mm^2. The fODFs are superimposed over the background of their related IDMs.

reliably across the entire brain, thereby avoiding the need to restrict the computations to the regions of white matter. In particular, in addition to reconstruction of fODFs $f_a(\mathbf{u}|\mathbf{r})$, the algorithm can also yield a useful estimation of its related IDM $f_{iso}(\mathbf{r})$, which quantifies a relative contribution of the isotropic diffusion component as well as its spatial pattern. Moreover, one of the principal contributions of the present work is to demonstrate the effectiveness of exploiting *different* prior models for regularization of the spatial-domain behaviour of the reconstructed fODFs and IDMs. Specifically, the fibre continuity model of [52] has been used to force the local maxima of the fODFs to vary consistently throughout the brain, whereas the bounded variation model of [59] has helped us to achieve piecewise smooth reconstructions of the IDMs (which appear to be in a good agreement with brain anatomy). Moreover, although the results of our experimental study support the above prior models, the latter are by no means exclusive, and hence additional improvements could still be achieved through the use of more advanced methods of spatial regularization.

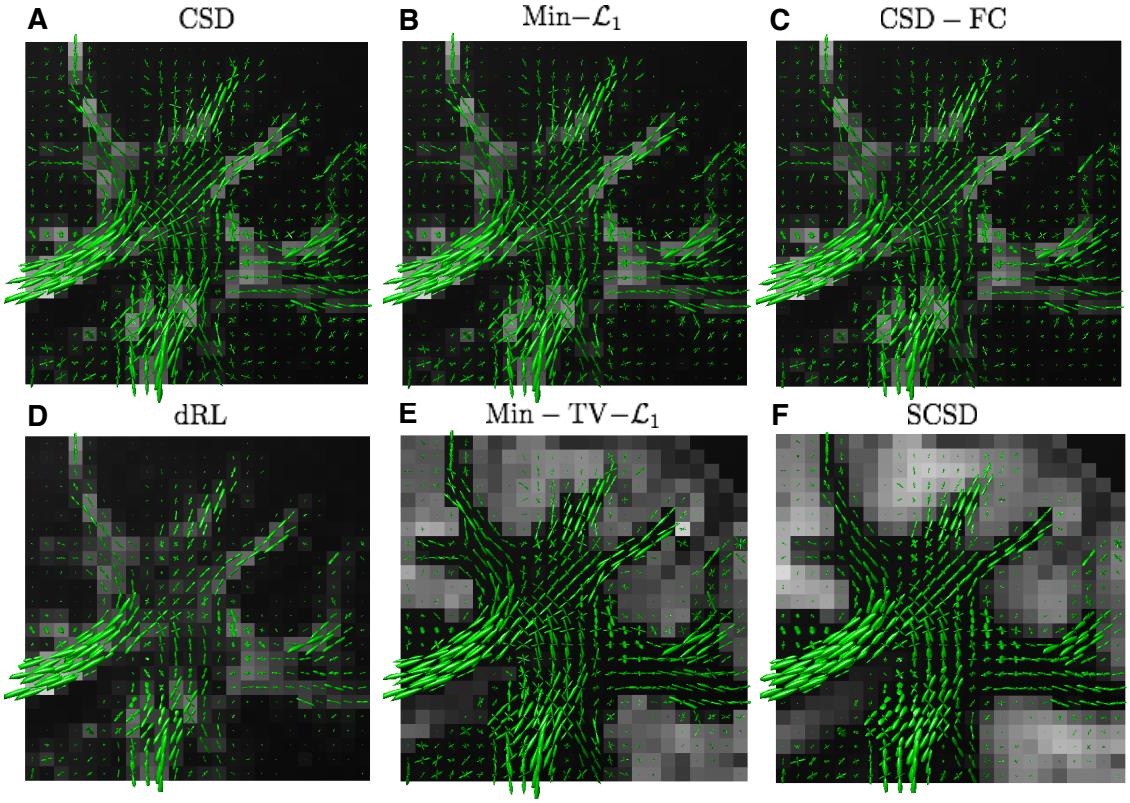

FIGURE 14. Estimated fODFs computed by CSD (A), Min-$\mathcal{L}_1$ (B), CSD-FC (C), dRL (D), Min-TV-$\mathcal{L}_1$ (E), and SCSD (F) for the coronal view and $b = 3000$ s/mm^2. The fODFs are superimposed over the background of their related IDMs.

Thus, for example, the bounded variation model could have been replaced by requiring the diffusion-encoded images $s_k(\mathbf{r})$ to have sparse representation coefficients in the domain of a multi-resolution transform (e.g., framelets [66]). At the same time, the linear filtering resulting from the fibre continuity assumption (as explained in the Appendix) could have been potentially replaced by a non-linear spatially-adaptive filtering. Whether or not the above modifications can result in substantial improvements in the quality of SD reconstruction, it remains a subject of our further research.

The estimated IDMs could serve as an additional input to fibre tractography or an independent observation for further statistical inference. In both scenarios, we believe the spatial regularization is essential to make the reconstructions of fODFs and IDMs be in a better agreement with the anatomical structure of the brain. Furthermore, as long as fibre tractography is concerned, our experimental results indicate the possibility to achieve a better angular resolution using higher values of b (e.g., $b = 3000$ s/mm^2). It is important

to note that this outcome would not have been possible without regularizing the IDMs (as explained in Section 4 and experimentally verified in Section 6) due to the adverse effect of measurement noises.

The proposed SCSD algorithm has been formulated as a convex minimization problem, which admits a unique and stable minimizer. Moreover, an important contribution of this work is to introduce a computationally efficient scheme for numerical implementation of SCSD. In particular, using ADMM, we have been able to find the optimal solution via a sequence of simpler optimization problems, which are both computationally efficient and amenable to parallel computations.

The performance of the SCSD method has been quantitatively compared against that of a number of alternative approaches, which have been specially selected to demonstrate the importance of various assumptions and constraints. Although SCSD has been shown to outperform the reference methods in terms of all the comparison metrics used, it still has a few limitations, whose mitigation could potentially produce more accurate reconstructions. Thus, from a Bayesian perspective, the cost functional in (15) suggests the measurement noise to be of an additive Gaussian nature, while in practice it is more likely to be Rician. It should be noted, however, that the Gaussian model can serve as a good approximation to the Rician model for relative high values of SNR (i.e., SNR $\gg 5$) [67]. Besides, using the Gaussian model is essential for rendering the optimization problem convex and computationally tractable. Yet, we note that, if using the Rician statistics is preferred, it should be straightforward to modify the SCSD algorithm as detailed in [67] for a similar setting. Finally, it should also be mentioned that the present version of SCSD assumes the diffusion outside white matter to be purely isotropic, while evidence exists that the diffusion within grey matter could exhibit a certain degree of anisotropy [68]. To account for this phenomenon, one could, for instance, replace the constant column of Φ in (7) by spherical harmonics up to second degree inclusive. Exploring such a modification and its effect on the quality of SD reconstruction defines another direction of our future research.

Appendix

To find a computationally efficient way to solve (26), it is convenient to slightly simplify the notations first. To this end, we note that the problem has a general form of

$$\min_{w} \left\{ \frac{1}{2}\|w - q\|_2^2 + \tau \|\mathbf{T}_{\mathbf{v}} w\|_2^2 \right\}, \tag{32}$$

with w interpreted as an I-dimensional (row) vector to be optimized over, q is a data vector, $\tau > 0$ is a fixed regularization constant, and $\mathbf{T}_{\mathbf{v}} : \mathbb{R}^I \to \mathbb{R}^I$ is the operator of *directional differencing* in the direction of $\mathbf{v} \in \mathbb{S}^2$. The problem (32) is a standard LS problem, whose optimal solution $w^\star$ satisfies a system of normal equations of the form

$$(\mathbf{I} + 2\tau\, \mathbf{T}_{\mathbf{v}}^* \mathbf{T}_{\mathbf{v}})\, w^\star = q, \tag{33}$$

with $\mathbf{T}_{\mathbf{v}}^*$ standing for the adjoint of $\mathbf{T}_{\mathbf{v}}$.

It goes without saying that a straightforward inversion of the matrix on the left-hand side of (33) cannot be accepted as a practical option, even for moderately sized diffusion-encoded images. Thus, a more efficient way to solve (33) needs to be found. To this end, we first note that, if the practical computation of partial differences was based on a standard backward-differencing scheme, then multiplication by $\mathbf{T}_{\mathbf{v}}$ in the spatial domain would be equivalent to linear filtering with frequency response $H_{\mathbf{v}}(\boldsymbol{\omega})$ given by

$$H_{\mathbf{v}}(\boldsymbol{\omega}) = 2\sum_{d=1}^{3} \mathbf{v}_d \sin\frac{\boldsymbol{\omega}_d}{2} e^{-\jmath\left(\frac{\boldsymbol{\omega}_d-\pi}{2}\right)}, \quad \boldsymbol{\omega} \in \mathbb{R}^3, \tag{34}$$

with $\mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ and $\boldsymbol{\omega} = (\boldsymbol{\omega}_1, \boldsymbol{\omega}_2, \boldsymbol{\omega}_3)$. In this case, by letting $\hat{w}^\star(\boldsymbol{\omega})$ and $\hat{q}(\boldsymbol{\omega})$ denote the discrete-space Fourier transform (DSFT) of w and q, respectively, the optimal solution in (33) could be defined in the DSFT domain as

$$\hat{w}^\star(\boldsymbol{\omega}) = \left(1 + 2\tau|H_{\mathbf{v}}(\boldsymbol{\omega})|^2\right)^{-1} \hat{q}(\boldsymbol{\omega}), \tag{35}$$

which suggests that $w^\star(\boldsymbol{\omega})$ is, in fact, a linearly filtered version of q.

Although it is definitely possible to use (35) in practical computations, such an approach would not be recommended for two main reasons. First, FFT-based implementation of linear filtering would entail the use of periodic boundary conditions, which might not be natural for the case at hand. Second, the logarithmic complexity of FFT might still be considered to be prohibitively high (especially taking into account the fact that the filter needs to be applied J times for J different values of $\mathbf{v}$).

A practical alternative to an FFT-based computation of $\hat{w}^\star(\boldsymbol{\omega})$ could be first to transform the frequency response $\left(1 + 2\tau|H_{\mathbf{v}}(\boldsymbol{\omega})|^2\right)^{-1}$ back into the spatial domain, followed by truncating the impulse response thus obtained. In our experiments, the impulse responses have been truncated to a size of $7 \times 7 \times 7$ voxels, with an associated error being less than 1% for a typical choice of the parameter τ (or, equivalently, μ). The resulting finite impulse response filters have distinctive (directional) low-pass characteristics, and they can be easily applied with a linear complexity under arbitrary boundary conditions.

References

[1] H. Johansen-Berg and T. E. J. Behrens. *Diffusion MRI: From quantitative measurements to in-vivo neuroanatomy*. Academic Press, first edition, 2009.

[2] P. J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophys. J.*, 66(1):259–267, 1994.

[3] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson. Imag., Ser. B*, 103(3):247–254, 1994.

[4] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurological disorders. *Radiology*, 161:401–407, 1986.

[5] D. Le Bihan, J.-F. Mangin, C. Poupon, C. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging*, 13:534–546, 2001.

[6] S. Mori. *Introduction to diffusion tensor imaging*. Elsevier, 2007.

[7] L. Frank. Anisotropy in high angular resolution diffusion-tensor MRI. *Magnetic Resonance in Medicine*, 45:935–939, 2001.

[8] L. R. Frank. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 47:1083–1099, 2002.

[9] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48:577–582, 2002.

[10] V. J. Wedeen, P. Hagmann, W.-Y. Tseng, T. G. Reese, and R. M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54:1377–1386, 2005.

[11] M. Descoteaux, R. Deriche, Denis LeBihan, J.-F. Mangin, and C. Poupon. Multiple q-shell diffusion propagator imaging. *Medical Image Analysis*, 15(4):603–621, August 2011.

[12] M. I. Menzel, E. T. Tan, K. Khare, J. I. Sperl, K. F. King, X. Tao, C. J. Hardy, and L. Marinelli. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. *Magnetic Resonance in Medicine*, 66:1226–1233, 2011.

[13] S. Merlet and D. Rachid. Compressed sensing for accelerated EAP recovery in diffusion MRI. In *Proceedings of MICCAI*, pages 14–21, Pekin, Chine, September 2010.

[14] Y.-C. Wu and A. L. Alexander. Hybrid diffusion imaging. *NeuroImage*, 36:617–629, 2007.

[15] J. H. Jensen, J. A. Helpern, A. Ramani, H. Lu, and K. Kaczynski. Diffusional kurtosis imaging: The quantification of non-Gaussian water diffusion by means of magnetic resonance imaging. *Magnetic Resonance in Medicine*, 53:1432–1440, 2005.

[16] Y. Assaf, R. Z. Freidlin, G. K. Rohde, and P. J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52:969–978, 2004.

[17] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander. NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage*, 61(4):1000–1016, 2012.

[18] H. Assemlal, D. Tschumperle, and L. Brun. Efficient computation of pdf-based characteristics from diffusion MR signal. In *Lecture Notes in Computer Science*, pages 70–78. MICCAI, 2008.

[19] S. Merlet, J. Cheng, A. Ghosh, and R. Deriche. Spherical polar Fourier EAP and ODF reconstruction via compressed sensing in diffusion MRI. In *Proceedings of ISBI*, pages 365–371, 2011.

[20] E. Ozarslan, C. G. Koay, T. M. Shepherd, M. E. Komlosh, M. O. Irfanoglu, and C. Pierpaoli P. J. Basser. Mean apparent propagator (MAP) MRI: A novel diffusion imaging method for mapping tissue microstructure. *NeuroImage*, 78:16–32, 2013.

[21] D. S. Tuch, T. G. Reese, M. R. Wiegell, and V. J. Wedeen. Diffusion MRI of complex neural architecture. *Neuron*, 40:885–895, 2003.

[22] D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–1372, 2004.

[23] D. C. Alexander. Multiple-fiber reconstruction algorithms for diffusion MRI. *Annals of the New York Academy of Science*, 1064:113–133, 2005.

[24] C. P. Hess, P. Mukherjee, E. T. Han, D. Xu, and D. R. Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine*, 56:104–117, 2006.

[25] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine*, 58:497–510, 2007.

[26] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Apparent diffusion coefficients from high angular resolution diffusion images: Estimation and applications. *Magnetic Resonance in Medicine*, 56(2):395–410, 2006.

[27] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel. Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magnetic Resonance in Medicine*, 64:554–566, 2010.

[28] A. Barnett. Theory of Q-ball imaging redux: Implications for fiber tracking. *Magnetic Resonance in Medicine*, 62:910–923, 2009.

[29] A. Tristan-Vega, C. F. Westin, and S. Aja-Fernandez. Estimation of fiber orientation probability density functions in high angular resolution diffusion imaging. *NeuroImage*, 47:638–650, 2009.

[30] D. Le Bihan, C. Poupon, A. Amadon, and F. Lethimonnier. Artifacts and pitfalls in diffusion MRI. *Journal of Magnetic Resonance Imaging*, 24:478–488, 2006.

[31] O. Pasternak, N. Sochen, Y. Gur, N. Intrator, and Y. Assaf. Free water elimination and mapping from diffusion MRI. *Magnetic Resonance in Medicine*, 62(3):717–730, 2009.

[32] J. Malcolm, M. Shenton, and Y. Rathi. Neural tractography using an unscented Kalman filter. In *Proceedings of IPMI*, pages 126–138, 2009.

[33] D. M. Healy Jr., H. Hendriks, and P. T. Kim. Spherical deconvolution. *J Multivar Anal*, 67:1–22, 1998.

[34] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23:1176–1185, 2004.

[35] J.-D. Tournier, F. Calamante, and A. Connelly. Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage*, 35:1459–1472, 2007.

[36] J.-D. Tournier, C.-H. Yeh, F. Calamante, K.-H. Cho, A. Connelly, and C.-P. Lin. Resolving crossing fibres using constrained spherical deconvolution: Validation using diffusion-weighted imaging phantom data. *NeuroImage*, 42:617–625, 2008.

[37] B. Jeurissen, A. Leemans, D. K. Jones, J.-D. Tournier, and J. Sijbers. Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution. *Human brain mapping*, 32(3):461–79, 2011.

[38] A. W. Anderson. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54:1194–1206, 2005.

[39] K. E. Sakaie and M. J. Lowe. An objective method for regularization of fiber orientation distributions derived from diffusion-weighted MRI. *NeuroImage*, 34:169–176, 2007.

[40] D. C. Alexander. Maximum entropy spherical deconvolution for diffusion MRI. In *Proceed. of IPMI*, pages 76–87, 2005.

[41] M. Jing, T. M. McGinnity, S. Coleman, H. Zhang, A. Fuchs, and J. A. S. Kelso. Enhancement of fiber orientation distribution reconstruction in diffusion-weighted imaging by single channel blind source separation. *IEEE Transactions on Medical Imaging*, 59(2):363–373, 2012.

[42] B. Jian and B. C. Vemuri. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Transactions on Medical Imaging*, 26(11):1464–1471, November 2007.

[43] F. Dell'Acqua, G. Rizzo, P. Scifo, R. A. Clarke, G. Scotti, and F. Fazio. A model-based deconvolution approach to solve fiber crossing in Diffusion-Weighted MR Imaging. *IEEE Trans. Biomed. Eng.*, 54(3), March 2007.

[44] F. Dell'Acqua, P. Scifo, G. Rizzo, M. Catani, A. Simmons, G. Scotti, and F. Fazio. A modified damped Richardson-Lucy algorithm to reduce isotropic background effects in spherical deconvolution. *NeuroImage*, 49:1446–1458, 2010.

[45] T. Schultz and H. Seidel. Estimating crossing fibers: A tensor decomposition approach. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1635–1642, 2008.

[46] P.-T. Yap and D. Shen. Spatial transformation of DWI data using non-negative sparse representation. *IEEE Transactions on Medical Imaging*, 31(11):2035–2049, November 2012.

[47] E. Kaden, A. Anwander, and T. R. Knosche. Variational inference of the fiber orientation density using diffusion mr imaging. *NeuroImage*, 42:1366–1380, 2008.

[48] E. Kaden and F. Kruggel. Nonparametric Bayesian inference of the fiber orientation distribution from diffusion-weighted MR images. *Medical Image Analysis*, 16:876–888, 2012.

[49] A. Goh, C. Lenglet, P. M. Thompson, and R. Vidal. Estimating orientation distribution functions with probability density constraints and spatial regularity. In *Lecture Notes in Computer Science*, volume 5761, pages 877–885. MICCAI, 2009.

[50] A. Raj, C. Hess, and P. Mukherjee. Spatial HARDI: Improved visualization of complex white matter architecture with Bayesian spatial regularization. *NeuroImage*, 54:396–409, 2011.

[51] A. Ramirez-Manzanares, M. Rivera, B. C. Vemuri, P. Carney, and T. Mareci. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *IEEE Transactions on Medical Imaging*, 26(8):1091–1102, August 2007.

[52] M. Reisert and V. G. Kiselev. Fiber continuity: An anisotropic prior for ODF estimation. *IEEE Transactions on Medical Imaging*, 30(6):1274–1283, June 2011.

[53] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, editors, *Springer Optimization and Its Applications*, volume 49. Springer, 2011.

[54] W. Freeden and F. Schreiner. Orthogonal and non-orthogonal multiresolution analysis, scale discrete and exact fully discrete wavelet transform on the sphere. *Constructive Approximation*, 14:493–515, 1998.

[55] G. D. Parker, D. Marshall, P. L. Rosin, N. Drage, S. Richmond, and D. K. Jones. A pitfall in the reconstruction of fibre ODFs using spherical deconvolution of diffusion MRI data. *NeuroImage*, 65:433–448, 2013.

[56] O. Michailovich and D. Adam. A high-resolution technique for ultrasound harmonic imaging using sparse representations in Gabor frames. *IEEE Transactions on Medical Imaging*, 21(12):1490–1503, December 2002.

[57] J. L. Starck, M. Elad, and D. Donoho. Redundant multiscale transforms and their application for morphological component analysis. *Adv. Imag. Electron. Phys.*, 132, 2004.

[58] I. Daubechies, M. Defrise, and C. DeMol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math*, 57(11):1413–1457, 2004.

[59] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4):259–268, 1992.

[60] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trend. Mach. Learn.*, 3(1), 2010.

[61] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[62] A. Chambolle and P.-L. Lions. Image recovery via total variation minimization and related problems. *Numer. Math.*, 76:167–188, 1997.

[63] E. B. Saff and A. B. J. Kuijlaars. Distributing many points on a sphere. *The Mathematical Intelligencer*, 19(1):5–11, December 1997.

[64] Q. Zhou, O. Michailovich, Y. Rathi, and A. Voineskos. Resolving complex fibre architecture by means of sparse spherical deconvolution in the presence of isotropic diffusion. In *Proceedings of SPIE Conference on Medical Imaging*, San-Diego, CA, 2014.

[65] J. P. Haldar and R. M. Leahy. Linear transforms for Fourier data on the sphere: Application to high angular resolution diffusion MRI of the brain. *NeuroImage*, 71:233–247, 2013.

[66] I. Daubechies, B. Han, Amos Ron, and Z. Shen. Framelets: MRA-based constructions of wavelet frames. *Applied and Computational Harmonic Analysis*, 14(1):1–46, Jan 2003.

[67] S. Dolui, I. C. Salgado Patarroyo, O. Michailovich, and Y. Rathi. Reconstruction of HARDI using compressed sensing and its application to contrast HARDI. In *Proceed. of MMBIA*, Breckenridge, CL, 2012.

[68] Y. Rathi, O. Pasternak, P. Savadjiev, O. Michailovich, S. Bouix, M. Kubicki, C.-F. Westin, N. Makris, and M. E. Shenton. Gray matter alterations in early aging: A diffusion magnetic resonance imaging study. *Human Brain Mapping (to appear)*, 2014.